arXiv:1609.07007v1 [stat.CO] 22 Sep 2016

Fast symmetric additive covariance smoothing

Jona Cederbaum, Fabian Scheipl, Sonja Greven*

Department of Statistics, Faculty of Mathematics, Computer Science and Statistics, LMU Munich, Germany

Abstract

We propose a fast bivariate smoothing approach for symmetric surfaces that has a wide range of applications. We show how it can be applied to estimate the covariance function in longitudinal data as well as multiple additive covariances in functional data with complex correlation structures. Our symmetric smoother can handle (possibly noisy) data sampled on a common, dense grid as well as irregularly or sparsely sampled data. Estimation is based on bivariate penalized spline smoothing using a mixed model representation and the symmetry is used to reduce computation time compared to the usual non-symmetric smoothers. We outline the application of our approach in functional principal component analysis and demonstrate its practical value in two applications. The approach is evaluated in extensive simulations. We provide documented open source software implementing our fast symmetric bivariate smoother building on established algorithms for additive models.

Key words: functional data, longitudinal data, functional principal components, penalized splines

*__Address for correspondence:__ Sonja Greven, Department of Statistics, Faculty of Mathematics, Computer Science and Statistics, LMU Munich, Ludwigstr. 33, 80539 Munich, Germany.
E-mail: Sonja.Greven@stat.uni-muenchen.de

1 Introduction

Covariance functions play a central role in many areas of statistics. They summarize the dependency between stochastic observations and encode smoothness assumptions about (observed or latent) random processes. We propose a fast bivariate smoothing approach for symmetric surfaces which can estimate covariance functions in a wide range of data situations. Our approach can handle dependent processes based on an additive decomposition of the covariance function and is also applicable to processes that are observed on irregular or sparse grids.

In functional data analysis (FDA; see, e.g. Ramsay and Silverman, 2005), covariance functions are at the heart of functional principal component analysis (FPCA), a key tool for dimension reduction based on an eigen analysis of the covariance operator of a random process. FPCA is commonly used to estimate the model parameters in functional predictor and functional response regression models (see Morris, 2015, for an overview). Other examples that are based on covariance functions include functional discriminant analysis (James and Hastie, 2001) and functional canonical correlation analysis (Leurgans et al., 1993). In longitudinal data analysis (LDA), where measurements are frequently recorded at irregularly spaced time points, the correct specification of the covariance benefits the estimation efficiency of the fixed effects and improves the individual predictions (cf. Fan et al., 2007). The covariance is also a crucial ingredient in time series analysis, e.g. in risk models and portfolio allocation (cf. Tai, 2009). The interest commonly lies in a single time series in contrast to FDA (and LDA) where multiple curves are observed, e.g. over time. In principle, our symmetric smoothing approach is also applicable to time series which is, however, not the focus in this paper.

Covariance functions are commonly assumed to be smooth. Thus, when the observed curves are not sufficiently smooth (i.e. observed with error) or not measured on a common dense grid, smoothing becomes necessary at some point during covariance estimation. Directly smoothing observed curves (see, e.g., Besse and Ramsay, 1986), however, is very difficult or impossible for sparsely observed data which are frequently recorded both in FDA and LDA (Yao et al., 2005). Moreover, pre-smoothing observed curves removes the measurement error, which is not accounted for in subsequent estimation steps. We pur-

sue an alternative approach and apply bivariate smoothing to the sample covariance of the observed data points.

Most existing work on non-parametric covariance estimation is either restricted to independent functional (or longitudinal) observations and/or only applies to data sampled on a common grid. Furthermore, most bivariate smoothing approaches are not specifically designed for covariances. They do not exploit the symmetry of the estimated surface and thus use redundant information in the available data. To the best of our knowledge, previous approaches have never addressed these issues simultaneously. They can be divided according to three main criteria: 1) the generality of the assumed correlation structure in the data, 2) the generality of possible sampling grids, and 3) the estimation procedure including the selection of the degree of smoothing.

A number of approaches address covariance smoothing in LDA. They are restricted to independent curves but allow for general sampling grids. Smoothing is either accomplished by bivariate kernel smoothing (e.g. Staniswalis and Lee, 1998; Yao et al., 2003, 2005) or by bivariate (penalized) spline smoothing (e.g. Kauermann and Wegener, 2011). The degree of smoothing is either chosen by visual inspection (Staniswalis and Lee, 1998), different leave-one-curve-out cross validation algorithms (e.g. Yao et al., 2003, 2005) or based on a mixed model representation (e.g. Kauermann and Wegener, 2011). These approaches do not account for the symmetry of the estimated surface. James et al. (2000) directly estimate the smooth eigenfunctions of the covariance function. They estimate a reduced rank mixed effects model via the EM algorithm and use B-spline basis functions to represent the eigenfunctions of the covariance operator. Peng and Paul (2009) estimate the same reduced rank model based on a more efficient Newton-Raphson procedure on the Stiefel manifold. The extension of these reduced rank methods to complex correlation structures is not straightforward. Xiao et al. (2016b) recently proposed a bivariate smoother designed for covariance smoothing which can be used for sparsely observed, independent functions. They use bivariate penalized B-splines and enforce a symmetry constraint on the spline coefficients which we take up in our extension to correlated curves. Estimation is done by a three-step procedure which accounts for the covariance of the sample covariance. Their leave-one-curve-out cross validation procedure for selecting the smoothing parameter is not applicable for correlated functional data, however.

Other covariance smoothing approaches can be applied to correlated functions but are restricted to functions sampled on a common grid and considerably simpler correlation structures than ours. Di et al. (2009) and Greven et al. (2010) use bivariate penalized splines and select the smoothing parameter using restricted maximum likelihood (REML; Patterson and Thompson, 1971) estimation. Shou et al. (2015) apply a method of moments approach based on symmetric sums represented in a sandwich form. For smoothing, they propose to use an extension of the fast covariance estimation algorithm of Xiao et al. (2016c) to correlated functions. Di et al. (2014) extend the functional random intercept model of Di et al. (2009) to sparsely sampled functional data, but the correlation structure remains less general than ours and an extension is not straightforward. More general correlation structures are allowed in the approach of Cederbaum et al. (2016) that is also suitable for sparsely and irregularly sampled functional data. Their focus lies, however, on a model with crossed functional random effects and estimation is only discussed for this special case. Apart from considering less general correlation structures, all these approaches neither avoid the use of redundant information nor account for the symmetry of the smoothed surface.

We propose a fast symmetric bivariate smoothing approach that applies to data with a broad range of possible correlation structures, much broader than existing methods. Furthermore, our approach is well-suited for (possibly noisy) data sampled on a common, dense grid as well as for irregularly or sparsely sampled data. Strength is borrowed by pooling information across different curves, which is particularly important for curves observed on sparse, unequal grids. The smoothing approach we present is widely applicable: In this paper, we demonstrate how it can be applied to longitudinal data as a special case of independent functional data as well as to correlated functional data with very general and complex correlation structures. For the latter, we extend our bivariate smoothing approach to smoothing additive covariance functions. To the best of our knowledge, all previous proposals in this field have been restricted to estimating much less general dependency structures.

We estimate the covariance functions using a smooth method of moments approach represented as a bivariate additive varying coefficient model. The estimation is based on bivariate penalized splines. We choose the smoothing parameters using REML, which allows the direct extension to additive bivariate smoothing of a superposition of multiple covariance

functions. This allows our method to be used for a broad range of complex real-word data settings. It also frees us from having to pre-specify a discrete grid of candidate values for the smoothing parameters that is required for cross-validation based approaches like Xiao et al. (2016b). Smoothing the sample covariance quickly becomes a high-dimensional problem as the number of elements in the sample covariance increases quadratically with the number of grid points. We take advantage of the symmetry of the sample covariance and only estimate the upper triangle of the surface including the diagonal. The estimates are then reflected across the diagonal to obtain the entire estimated covariance, which is continuous but not necessarily smooth across the diagonal. To avoid boundary effects on the diagonal, we enforce smoothness by imposing a symmetry constraint on the spline coefficients, which for the simplest case of independent curves reduces to that of Xiao et al. (2016b). We show how the symmetry constraint can be applied separately to additive covariances and can even be used for any bivariate symmetric smoothing problem beyond covariance functions. Our approach modifies the covariance smoothing approach proposed in Cederbaum et al. (2016) and extends it to more general models. It reduces both the data entering the estimation and the number of spline coefficients that have to be estimated, which leads to considerably faster estimation requiring less memory.

We provide software implementing our approach based on a novel constructor function for R-package mgcv, which provides a general framework for additive models allowing for a very flexible model specification (R Development Core Team, 2014; Wood, 2011).

We outline the application of our approach to FPCA and demonstrate its practical relevance by an application to sparse longitudinal observations of CD4 cell count trajectories and to densely but irregularly observed acoustic signals from a speech production study. This study requires crossed functional random effects due to repeated measurements for both speakers and target words and thus corresponds to a case of dependent functional data with additive covariance structure.

We organize our paper as follows: Section 2 first develops our fast symmetric covariance smoothing approach for a simple special case with only one smooth covariance function and additional measurement error. In Section 3, the smoother is extended to complex dependency structures involving the smoothing of multiple additive covariance functions. Section 4 outlines the application of our covariance smoother in FPCA. In Section 5, details

on the implementation are given. In Section 6 and 7, we evaluate our approach in an application to speech production data and in simulations, respectively. Section 8 closes with a discussion and outlook. Theoretical results and supplementary material are available in the appendix. R-code implementing our approach can be provided upon request.

2 Fast symmetric covariance smoothing

For simplicity, we first explain our covariance smoothing approach for a simple special case with only one smooth auto-covariance and additional measurement error. This will be extended to additive covariance smoothing in Section 3.

2.1 Model with independent curves

Consider the following model

$$Y_i(t_{ij}) = \mu(t_{ij}, \boldsymbol{x}_i) + E_i(t_{ij}) + \varepsilon_i(t_{ij}),\ j = 1, \dots, D_i,\ i = 1, \dots, n, \tag{2.1}$$

where $Y_i(t_{ij})$ is the observed value of response curve i observed at the jth observation point $t_{ij} \in \mathcal{T}$, a closed interval in $\mathbb{R}$. $\mu(t_{ij}, \boldsymbol{x}_i)$ is a global mean function depending on a vector of known covariates $\boldsymbol{x}_i$. $E_i(t_{ij})$ is a smooth curve-specific deviation from the global mean and $\varepsilon_i(t_{ij})$ is additional independent and identically distributed white noise measurement error with constant variance σ^2 that accounts for random uncorrelated variation within curve i. The model can be seen as a function-on-scalar regression model (e.g. Faraway, 1997; Ramsay and Silverman, 2005; Reiss et al., 2010) where all n curves are assumed to be independent and it is a special case of the general functional linear mixed model (FLMM; see Morris, 2015, for a discussion and further references) in Section 3 with one curve-specific functional random intercept. Model (2.1) is often applied to longitudinal data with $\mathcal{T}$ denoting a time interval. Note that all curves may either be observed on a common, fine grid or on curve-specific, possibly sparse, D_i evaluation points t_{ij}, $j = 1, \dots, D_i$, $i = 1, \dots, n$.

In the following, we assume that $E_i(\cdot)$ and $\varepsilon_i(\cdot)$, $i = 1, \dots, n$, are zero-mean, mutually uncorrelated random processes and that $E_i(\cdot)$ is square-integrable. We denote the auto-

covariance function of $E_i(\cdot)$ by $K^E(t,t') = \text{Cov}\left[E_i(t), E_i(t')\right]$, $t, t' \in \mathcal{T}$. We further assume that the mean and the auto-covariance are smooth in t and in arguments t, t', respectively.

2.2 Estimation in the independent case

We apply the following smooth method of moments approach to estimate the auto-covariance $K^E(t,t')$. It modifies the approach presented in Cederbaum et al. (2016) by accounting for the symmetry of covariances, which leads to a considerable reduction of computation times. In this simple case of independent functional responses, our smoother is closely related to that of Xiao et al. (2016b), who approach the problem from a slightly different perspective. Their approach, however, is not directly extendable to correlated data as will be discussed in Section 3.

We focus in the following on the centered functional responses $\tilde{Y}_i(t_{ij}) = Y_i(t_{ij}) - \mu(t_{ij}, \boldsymbol{x}_i)$ with expectation zero and denote their realizations by $\tilde{y}_i(t_{ij})$. We exploit the fact that the expectation of the centered cross products corresponds to the covariance of the functional response which is given as

$$\mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_i(t_{ij'})\right] = \text{Cov}\left[\tilde{Y}_i(t_{ij}), \tilde{Y}_i(t_{ij'})\right] = K^E(t_{ij}, t_{ij'}) + \sigma^2\delta_{jj'}, \qquad (2.2)$$
$$j, j' = 1, \ldots, D_i,\ i = 1, \ldots, n,$$

with $\delta_{jj'}$ equal to one if $j = j'$ and zero otherwise. Equation (2.2) can be seen as a special case of a bivariate additive varying coefficient model for the empirical covariances $\tilde{y}_{it_{ij}}\tilde{y}_{it_{ij'}}$, in which the auto-covariance and the error variance are the unknown components. We estimate the smooth auto-covariance function $K^E(t,t')$ and the error variance σ^2 simultaneously using a bivariate spline representation for $K^E(t,t')$ under working assumptions of independence and homoscedasticity.

For this, let $\boldsymbol{C}$ denote the $\mathcal{C} \times 1$ stacked vector of all cross products, with $\mathcal{C} = \sum_{i=1}^n D_i^2$. Then, Model (2.2) can be represented as

$$\mathbb{E}(\boldsymbol{C}) \;=\; \left[\boldsymbol{M}^E | \boldsymbol{\delta}^\varepsilon\right] \left(\boldsymbol{\theta}^{E^\top}, \sigma^2\right)^\top =: \boldsymbol{M}\boldsymbol{\alpha}, \qquad (2.3)$$

where $\boldsymbol{M}^E$ denotes the $\mathcal{C} \times \left(F^E\right)^2$ bivariate spline design matrix, containing the evaluations of any bivariate spline basis with $\left(F^E\right)^2$ basis functions that are symmetric across the diagonal, i.e. across $t_{ij} = t_{ij'}$. We use bivariate tensor product B-splines, but other bases are possible. See Section 5 and Appendix B for details. $\boldsymbol{\delta}^{\varepsilon}$ is an indicator vector of length $\mathcal{C}$ whose elements take values $\delta_{jj'}$. $\boldsymbol{\theta}^E$ is a spline coefficient vector of length $\left(F^E\right)^2$. To avoid over-fitting, we use an isotropic quadratic smoothness penalty of the form

$$\text{pen}(\lambda) = \lambda {\boldsymbol{\theta}^E}^\top \boldsymbol{S}^E \boldsymbol{\theta}^E, \tag{2.4}$$

where λ denotes the smoothing parameter that controls the bias-variance tradeoff and $\boldsymbol{S}^E$ is a suitable penalty matrix, see Section 5 for a discussion.

The development above uses all $\mathcal{C}$ cross products for the estimation of the covariance as is commonly done in published smooth method of moments approaches for covariance estimation (Yao et al., 2003; Yao et al., 2005; for correlated curves, e.g., Staniswalis and Lee, 1998; Di et al., 2009; Greven et al., 2010). As it is quadratic in the number of function evaluations, $\mathcal{C}$ quickly becomes extremely large in practice and often poses significant computational challenges. Since the cross products $\tilde{y}_{it_{ij}}\tilde{y}_{it_{ij'}}$ and $\tilde{y}_{it_{ij'}}\tilde{y}_{it_{ij}}$ are identical, we, like Xiao et al. (2016b), avoid the use of this redundant information and only estimate the upper triangle of the auto-covariance surface including the diagonal. Moreover, using this redundant data twice violates the implicit working assumption of independent observations in our additive model. A detailed discussion of working assumptions and possible strategies for handling violations is given in Section 3.4.

We assume in the following that $\boldsymbol{C}$ is sorted such that it can be partitioned as $\boldsymbol{C} = \left({\boldsymbol{C}_{t<t'}}^\top, {\boldsymbol{C}_{t=t'}}^\top, {\boldsymbol{C}_{t>t'}}^\top\right)^\top$, where $\boldsymbol{C}_{t<t'}$, $\boldsymbol{C}_{t=t'}$, $\boldsymbol{C}_{t>t'}$ comprise all cross products $\tilde{y}_{it_{ij}}\tilde{y}_{it_{ij'}}$, $i = 1, \ldots, n$, with $t_{ij} < t_{ij'}$, $t_{ij} = t_{ij'}$, and $t_{ij} > t_{ij'}$, respectively. Then, with suitable sorting within the three partitions of $\boldsymbol{C}$, the symmetry of the cross products implies that $\boldsymbol{C}_{t<t'} = \boldsymbol{C}_{t>t'}$. In order to speed up estimation, we only use $\boldsymbol{C}^{\Delta} \coloneqq \left({\boldsymbol{C}_{t<t'}}^\top, {\boldsymbol{C}_{t=t'}}^\top\right)^\top$ for the estimation in the bivariate model (2.3). The total number of cross products thus amounts to $\mathcal{C}^{\Delta} := \sum_{i=1}^{n} D_i(D_i+1)/2$. The design matrix of the bivariate additive model can accordingly be partitioned as $\boldsymbol{M} = \left({\boldsymbol{M}_{t<t'}}^\top, {\boldsymbol{M}_{t=t'}}^\top, {\boldsymbol{M}_{t>t'}}^\top\right)^\top$. Let $\boldsymbol{M}^{E\Delta}$ and $\boldsymbol{\delta}^{\epsilon\Delta}$ denote the submatrix and subvector corresponding to $\boldsymbol{C}^{\Delta}$. Then, the bivariate additive

model (2.3) reduces to $\mathbb{E}\left(\boldsymbol{C}^{\Delta}\right) = \left[\boldsymbol{M}^{E\Delta}|\boldsymbol{\delta}^{\varepsilon\Delta}\right]\left(\boldsymbol{\theta}^{E\top}, \sigma^2\right)^{\top}$. Reflecting the upper triangle across the diagonal ensures that the obtained surface estimate is symmetric and continuous in both directions t and t', but it does not guarantee smoothness across the diagonal without additional constraint. Boundary effects can occur because the coefficients can vary more freely, the smaller the overlap of the support of their associated basis functions with the area where $t \leq t'$, since less data are available for their estimation. This may lead to wiggly estimates in the area of the diagonal. An expected consequence is that the separation of the smooth auto-covariance surface and the error variance on the diagonal becomes more challenging, which we indeed observe in our application in Section 6 and in our simulations in Sections 7.

We enforce smoothness across the diagonal in order to avoid such boundary effects. As a symmetric surface implies a symmetric spline coefficient matrix $\boldsymbol{\Theta}^E = \left[\theta^E_{bb'}\right]_{b,b'=1,\ldots,F^E} = \boldsymbol{\Theta}^{E\top}$, where $\boldsymbol{\theta}^E = \left(\boldsymbol{\theta}^{E\top}_{b<b'}, \boldsymbol{\theta}^{E\top}_{b=b'}, \boldsymbol{\theta}^{E\top}_{b>b'}\right)^{\top}$ contains first the entries of $\boldsymbol{\Theta}^E$ below the diagonal ($\theta^E_{bb'}$, $b < b'$), then the diagonal entries ($\theta^E_{bb'}$, $b = b'$) and lastly the entries above the diagonal ($\theta^E_{bb'}$, $b > b'$), we impose a symmetry constraint on $\boldsymbol{\Theta}^E$. Thus, our approach differs in two crucial points from the naive covariance estimation in (2.3) and from most previous covariance smoothing approaches. First, we reduce the number of cross products that enter the estimation and second, imposing the symmetry constraint almost halves the number of spline coefficients that have to be estimated. Both aspects greatly speed up the computation as we show in Sections 6 and 7.

With suitable sorting within the partitions $\boldsymbol{\theta}^E_{b<b'}$ and $\boldsymbol{\theta}^E_{b>b'}$, the above symmetry constraint on the coefficient matrix corresponds to the following symmetry constraint on the coefficient vector

$$\theta^E_{bb'} \;=\; \theta^E_{b'b},\ b, b' = 1, \ldots, F^E \;\Leftrightarrow \boldsymbol{\theta}^E_{b<b'} = \boldsymbol{\theta}^E_{b>b'}, \tag{2.5}$$

which corresponds to the constraint used in Xiao et al. (2016b). This allows us to consider the reduced coefficient vector $\boldsymbol{\theta}^{Er} = \left(\boldsymbol{\theta}^{E\top}_{b<b'}, \boldsymbol{\theta}^{E\top}_{b=b'}\right)^{\top}$ of length $F^E\left(F^E+1\right)/2$. For the implementation of the above symmetry constraint, we use that under the constraint we have $\boldsymbol{M}^{E\Delta}\boldsymbol{\theta}^E = \left(\boldsymbol{M}^{E\Delta}_{b<b'} + \boldsymbol{M}^{E\Delta}_{b>b'}\right)\boldsymbol{\theta}^E_{b<b'} + \boldsymbol{M}^{E\Delta}_{b=b'}\boldsymbol{\theta}^E_{b=b'}$, with $\boldsymbol{M}^{E\Delta}_{b<b'}$, $\boldsymbol{M}^{E\Delta}_{b=b'}$, and $\boldsymbol{M}^{E\Delta}_{b>b'}$ containing

the respective columns of $\boldsymbol{M}^{E\Delta}$. Thus, the constraint is equivalent to adding up columns $\boldsymbol{M}^{E\Delta}_{b<b'}$ and $\boldsymbol{M}^{E\Delta}_{b>b'}$ of the design matrix. This can be achieved by right multiplication of $\boldsymbol{M}^{E\Delta}$ with the $\left(F^E\right)^2 \times F^E\left(F^E+1\right)/2$ constraint matrix

$$\boldsymbol{W}^E = \begin{bmatrix} \boldsymbol{I}_{\frac{F^E(F^E-1)}{2}} & \boldsymbol{0}_{\frac{F^E(F^E-1)}{2}\times F^E} \\ \boldsymbol{0}_{F^E\times\frac{F^E(F^E-1)}{2}} & \boldsymbol{I}_{F^E} \\ \boldsymbol{I}_{\frac{F^E(F^E-1)}{2}} & \boldsymbol{0}_{\frac{F^E(F^E-1)}{2}\times F^E} \end{bmatrix},$$

where $\boldsymbol{I}_x$ is an identity matrix of dimension x and $\boldsymbol{0}_{x\times y}$ is a null matrix of dimension $x\times y$. We denote the reduced design matrix by $\boldsymbol{M}^{E\Delta r} := \boldsymbol{M}^{E\Delta}\boldsymbol{W}^E$. Under the symmetry constraint, $\boldsymbol{M}^{E\Delta r}\boldsymbol{\theta}^{Er} = \boldsymbol{M}^{E\Delta}\boldsymbol{\theta}^E$. The penalty matrix in (2.4) also needs to be adjusted to the reduced coefficient vector $\boldsymbol{\theta}^{Er}$ as $\boldsymbol{S}^{Er} := \boldsymbol{W}^{E\top}\boldsymbol{S}^E\boldsymbol{W}^E$, corresponding to $\boldsymbol{\theta}^{Er\top}\boldsymbol{S}^{Er}\boldsymbol{\theta}^{Er} = \boldsymbol{\theta}^{E\top}\boldsymbol{S}^E\boldsymbol{\theta}^E$.

Altogether, the bivariate additive model for the reduced response vector with symmetry constraint (2.5) is given by

$$\mathbb{E}\left(\boldsymbol{C}^\Delta\right) \;=\; \left[\boldsymbol{M}^{E\Delta r}|\boldsymbol{\delta}^{\varepsilon\Delta}\right]\left(\boldsymbol{\theta}^{Er\top},\sigma^2\right)^\top =: \boldsymbol{M}^{\Delta r}\boldsymbol{\alpha}^r. \tag{2.6}$$

Note that in Model (2.6), each product $\tilde{y}_{it_{ij}}\tilde{y}_{it_{ij'}}$, $t_{ij}\le t_{ij'}$, enters the estimation with the same weight. This is not the case when all products $\tilde{y}_{it_{ij}}\tilde{y}_{it_{ij'}}$, $j,j'=1,\ldots,D_i$, are used as in Model (2.3), where all products appear twice except for those on the diagonal ($t_{ij}=t_{ij'}$). Our implementation allows to estimate Model (2.6) with the same weights as in Model (2.3) by putting a weight of 0.5 on the products on the diagonal. There is room for debate on whether it is desirable to down-weigh the data on the diagonal compared to the rest. One would expect that this leads to wigglier estimates but our simulations in Section 7 show that the difference is not very large.

In contrast to Xiao et al. (2016b), who derive a leave-one-curve-out generalized cross-validation (GCV) algorithm to choose the smoothing parameter for independent curves, we choose the smoothing parameter as variance component ratio using REML. REML has been shown to be more stable than GCV (Wood, 2011) and to be more robust to error correlation misspecification than prediction error methods (Krivobokova and Kauermann, 2007). Even more importantly, it allows us to directly extend our symmetric smoothing

approach to additive smoothing needed for functional data with complex dependency structures as will be shown in Section 3. In more general designs, where the responses cannot be decomposed into independent subvectors, it is not clear how to perform smoothing parameter selection based on GCV and optimizing multiple smoothing parameters would require a computationally costly multi-dimensional grid search.

Details on the implementation are given in Section 5 and in Appendix B.

3 Fast symmetric additive covariance smoothing

Simultaneous REML estimation of multiple smoothing parameters allows direct extension of our approach to more general models with complex correlation structures, for which we derive appropriate symmetry constraint matrices.

3.1 General functional linear mixed model

The general FLMM (see, e.g., Morris, 2015) can be seen as the functional analogue to the linear mixed model (LMM; see, e.g., Pinheiro and Bates, 2000), which is often applied to scalar correlated data. The random effects in the linear mixed model are replaced by functional random effects in order to account for the functional nature of the response. A functional random intercept (fRI) for a subject, for example, is a subject-specific deviation from the mean in form of a function. The FLMM is given by

$$Y_i(t_{ij}) = \mu(t_{ij}, \boldsymbol{x}_i) + \boldsymbol{z}_i^\top \boldsymbol{U}(t_{ij}) + E_i(t_{ij}) + \varepsilon_i(t_{ij}),\ j = 1, \ldots, D_i,\ i = 1, \ldots, n, \tag{3.1}$$

where $Y_i(t_{ij})$ denotes the response of curve i at observation point t_{ij}, which can be additively decomposed as in Model (2.1). Model (3.1), however, additionally accounts for correlation between (groups of) curves by the vector-valued random process $\boldsymbol{U}(t_{ij})$ which is multiplied by $\boldsymbol{z}_i$, a known covariate vector of length q. Examples for $\boldsymbol{z}_i^\top \boldsymbol{U}(t_{ij})$ yielding FLMMs with e.g. crossed and hierarchical functional random effects are given in Section 3.3 and in Appendix B.

We assume that $\boldsymbol{U}(\cdot)$, $E_i(\cdot)$, and $\varepsilon_i(\cdot)$ are zero mean, mutually uncorrelated random

processes and that $\boldsymbol{U}(\cdot)$ and $E_i(\cdot)$ are square-integrable. As for Model (2.1), we denote the auto-covariance of $E_i(\cdot)$ by $K^E(t,t') = \text{Cov}\left[E_i(t), E_i(t')\right]$, $t, t' \in \mathcal{T}$. The $q \times q$ matrix-valued auto-covariance of $\boldsymbol{U}(\cdot)$ is denoted by $\boldsymbol{K}^U(t,t') = \text{Cov}\left[\boldsymbol{U}(t), \boldsymbol{U}(t')\right]$. The covariances are assumed to be smooth (for each component in the case of $\boldsymbol{U}(t)$).

Let G denote the number of grouping variables. Then, $\boldsymbol{U}(t_{ij})$ can be divided into G independent blocks $\boldsymbol{U}_g(t_{ij})$, $g = 1, \ldots, G$, which again contain blocks of L^{U_g} independent copies $\boldsymbol{U}_{gl}(t_{ij})$, $l = 1, \ldots, L^{U_g}$, where L^{U_g} is the number of levels of the gth grouping variable. $\boldsymbol{U}_{gl}(t_{ij}) = \left(U_{gl1}(t_{ij}), \ldots, U_{gl\rho^{U_g}}(t_{ij})\right)^\top$, in turn, is a vector-valued random process of ρ^{U_g} components for each level of this grouping variable, for example $\rho^{U_g} = 2$ if the gth grouping variable is associated with a fRI and a functional random slope. The total number of entries in $\boldsymbol{U}(t_{ij})$ is given by $q = \sum_{g=1}^{G} L^{U_g}\rho^{U_g}$. The $\rho^{U_g} \times \rho^{U_g}$ matrix-valued covariance of $\boldsymbol{U}_{gl}(\cdot)$, $\boldsymbol{K}^{U_g}(t,t') = \left[K^{U_g}_{ss'}(t,t')\right]_{s,s'=1,\ldots,\rho^{U_g}} = \text{Cov}\left[\boldsymbol{U}_{gl}(t), \boldsymbol{U}_{gl}(t')\right]$, with $K^{U_g}_{ss'}(t,t') = K^{U_g}_{s's}(t',t)$, is the same for all levels, $l = 1, \ldots, L^{U_g}$, of the gth grouping variable. We can thus write the block-diagonal auto-covariance of $\boldsymbol{U}(\cdot)$ as

$$\boldsymbol{K}^U(t,t') = \text{diag}\left(\underbrace{\boldsymbol{K}^{U_1}(t,t'), \ldots, \boldsymbol{K}^{U_1}(t,t')}_{L^{U_1} \text{ times}}, \ldots, \underbrace{\boldsymbol{K}^{U_G}(t,t'), \ldots, \boldsymbol{K}^{U_G}(t,t')}_{L^{U_G} \text{ times}}\right).$$

3.2 Estimation in the general functional linear mixed model

Our fast symmetric covariance smoothing approach can be extended to the general model (3.1) by generalizing it to a matrix of covariances as described above and applying it to each additive component separately.

In analogy to Model (2.1), we base the covariance smoothing on the following decomposition of the expectation of the cross products of the centered functional responses

$$\begin{aligned} \mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\right] &= \text{Cov}\left[\tilde{Y}_i(t_{ij}), \tilde{Y}_{i'}(t_{i'j'})\right] \\ &= \boldsymbol{z}_i^\top \boldsymbol{K}^U(t_{ij}, t_{i'j'})\boldsymbol{z}_{i'} + \left[K^E(t_{ij}, t_{i'j'}) + \sigma^2\delta_{jj'}\right]\delta_{ii'}, \end{aligned} \tag{3.2}$$

where, in contrast to Model (2.1), products are now also computed across different curves i, i'.

Let $U_g(i)$ denote the level of U_g for observation i. Similar to $\boldsymbol{U}(\cdot)$, the covariate vector $\boldsymbol{z}_i$ can be divided into G blocks $\boldsymbol{z}_i^\top = \left(\boldsymbol{z}_i^{U_1\top}, \ldots, \boldsymbol{z}_i^{U_G\top}\right)$, where the blocks $\boldsymbol{z}_i^{U_g}$, $g = 1, \ldots, G$, can again be written as $\boldsymbol{z}_i^{U_g\top} = \left(\boldsymbol{z}_{i1}^{U_g\top}, \ldots, \boldsymbol{z}_{iL^{U_g}}^{U_g\top}\right)$ with $\boldsymbol{z}_{il}^{U_g\top} = \left(z_{il1}^{U_g}, \ldots, z_{il\rho^{U_g}}^{U_g}\right)$, $l = 1, \ldots, L^{U_g}$. The scalars $z_{ils}^{U_g}$ take the value of the respective covariate $\omega_{is}^{U_g}$ times an indicator $\delta_{U_g(i)l}$, specifying whether observation i belongs to level l of grouping variable g. Based on this partition, the expectation in (3.2) can be rewritten as

$$\begin{aligned}\mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\right] &= \sum_{g=1}^{G}\sum_{l=1}^{L^{U_g}}\sum_{s=1}^{\rho^{U_g}}\sum_{s'=1}^{\rho^{U_g}} z_{ils}^{U_g} z_{i'ls'}^{U_g} K_{ss'}^{U_g}(t_{ij}, t_{i'j'}) \\ &+ \left[K^E(t_{ij}, t_{i'j'}) + \sigma^2\delta_{jj'}\right]\delta_{ii'}.\end{aligned} \tag{3.3}$$

For example in the case of an FLMM with only one fRI ($G = 1$, $\rho^{U_1} = 1$), the products $z_{ils}^{U_g} z_{i'ls'}^{U_g}$ are indicators for whether the two observations in a cross products belong to the same level of the grouping variable.

We exploit the symmetry of covariances $\boldsymbol{K}^{U_g}(t_{ij}, t_{i'j'}) = \boldsymbol{K}^{U_g}(t_{i'j'}, t_{ij})^\top$, $g = 1, \ldots, G$, and of $K^E(t_{ij}, t_{i'j'}) = K^E(t_{i'j'}, t_{ij})$ and only use the products $\tilde{y}_{it_{ij}}\tilde{y}_{i't_{i'j'}}$ with $t_{ij} \leq t_{i'j'}$, suitably sorted in the long vector $\boldsymbol{C}^\Delta$. Let $\cdot$ denote the Hadamard (pointwise) product. As in the case with independent curves, Model (3.3) can be represented as bivariate additive varying coefficient model, here of the form

$$\mathbb{E}\left(\boldsymbol{C}^\Delta\right) = \left[\boldsymbol{M}^{U_1\Delta}|\ldots|\boldsymbol{M}^{U_G\Delta}|\boldsymbol{M}^{E\Delta}|\boldsymbol{\delta}^{\varepsilon\Delta}\right]\left(\boldsymbol{\theta}^{U_1\top}, \ldots, \boldsymbol{\theta}^{U_G\top}, \boldsymbol{\theta}^{E\top}, \sigma^2\right)^\top, \tag{3.4}$$

where $\boldsymbol{M}^{U_g\Delta}$, $g = 1, \ldots, G$, contain the column-wise concatenated submatrices $\boldsymbol{M}_{ss'}^{U_g\Delta}$, corresponding to the covariances $K_{ss'}^{U_g}(t, t')$, $s, s' = 1, \ldots, \rho^{U_g}$. The submatrices $\boldsymbol{M}_{ss'}^{U_g\Delta}$ are given by $\boldsymbol{Q}_{ss'}^{U_g\Delta} \cdot \boldsymbol{B}_{ss'}^{U_g\Delta}$, where $\boldsymbol{Q}_{ss'}^{U_g\Delta}$ contain suitably sorted and repeated entries $\delta_{U_g(i)U_g(i')} \cdot \omega_{is}^{U_g}\omega_{i's'}^{U_g}$ and $\boldsymbol{B}_{ss'}^{U_g\Delta}$ denote the bivariate spline design matrices. $\boldsymbol{M}^{E\Delta}$ is analogously given by $\boldsymbol{Q}^{E\Delta} \cdot \boldsymbol{B}^{E\Delta}$, with bivariate spline design matrix $\boldsymbol{B}^{E\Delta}$ and $\boldsymbol{Q}^{E\Delta}$ a new indicator matrix, which reduces to an all-ones matrix in the model with independent curves, for which thus $\boldsymbol{M}^{E\Delta} = \boldsymbol{B}^{E\Delta}$. The concrete form of $\boldsymbol{Q}_{ss'}^{U_g\Delta}$ and $\boldsymbol{Q}^{E\Delta}$, as well as the bivariate spline design matrices for tensor product B-splines are provided in Appendix B.

We assume that for each, $g = 1, \ldots, G$, the coefficient vector $\boldsymbol{\theta}^{U_g}$ is sorted correspond-

ingly to the columns of $\boldsymbol{M}^{U_g\Delta} = \left[\boldsymbol{M}_{11}^{U_g\Delta}|\ldots|\boldsymbol{M}_{1\rho^{U_g}}^{U_g\Delta}|\ldots|\boldsymbol{M}_{\rho^{U_g}1}^{U_g\Delta}|\ldots|\boldsymbol{M}_{\rho^{U_g}\rho^{U_g}}^{U_g\Delta}\right]$. Moreover, with suitable sorting, each submatrix $\boldsymbol{M}_{ss'}^{U_g\Delta}$ can be partitioned as in the case of independent curves, $\boldsymbol{M}_{ss'}^{U_g\Delta} = \left[\boldsymbol{M}_{ss',b<b'}^{U_g\Delta}|\boldsymbol{M}_{ss',b=b'}^{U_g\Delta}|\boldsymbol{M}_{ss',b>b'}^{U_g\Delta}\right]$. Let $\boldsymbol{\Theta}_{ss'}^{U_g}$ denote the correspondingly sorted coefficient matrices, where $\boldsymbol{\theta}_{ss'}^{U_g} = \left(\boldsymbol{\theta}_{ss',b<b'}^{U_g}{}^\top, \boldsymbol{\theta}_{ss',b=b'}^{U_g}{}^\top, \boldsymbol{\theta}_{ss',b>b'}^{U_g}{}^\top\right)^\top$ contains first the entries of $\boldsymbol{\Theta}^E$ below the diagonal, then the diagonal entries and lastly the entries above the diagonal. Assume further that within the three blocks of $\boldsymbol{\theta}_{ss'}^{U_g}$, the entries $\theta_{ss',bb'}$ are sorted correspondingly for all $s, s' = 1, \ldots, \rho^{U_g}$.

As a modular component, the symmetry constraint

$$\boldsymbol{\Theta}_{ss'}^{U_g} = \boldsymbol{\Theta}_{s's}^{U_g}{}^\top, \ s, s' = 1, \ldots, \rho^{U_g}, \tag{3.5}$$

can be applied to each, $g = 1, \ldots, G$, due to the symmetry of covariances $\boldsymbol{K}^{U_g}(t, t')$, yielding the reduced coefficient vectors $\boldsymbol{\theta}_{ss'}^{U_g r}$, $s \leq s'$, and thus the reduced long coefficient vector $\boldsymbol{\theta}^{U_g r}$. As in the case of independent curves, the constraint (3.5) is equivalent to adding up the respective columns of the large design matrix $\boldsymbol{M}^{U_g\Delta}$. This can be achieved by right-multiplication of $\boldsymbol{M}^{U_g\Delta}$ with a suitable constraint matrix $\boldsymbol{W}^{U_g}$, yielding the reduced design matrix $\boldsymbol{M}^{U_g\Delta r} = \left[\boldsymbol{M}_{11}^{U_g\Delta r}|\ldots|\boldsymbol{M}_{1\rho^{U_g}}^{U_g\Delta r}|\boldsymbol{M}_{22}^{U_g\Delta r}|\ldots|\boldsymbol{M}_{\rho^{U_g}-1\rho^{U_g}-1}^{U_g\Delta r}|\boldsymbol{M}_{\rho^{U_g}-1\rho^{U_g}}^{U_g\Delta r}|\boldsymbol{M}_{\rho^{U_g}\rho^{U_g}}^{U_g\Delta r}\right]$. Each $\boldsymbol{M}_{ss'}^{U_g\Delta r}$, $s \leq s'$, consists of column-wise concatenated matrices

$$\boldsymbol{M}_{ss'}^{U_g\Delta r} = \left[\boldsymbol{M}_{ss',b<b'}^{U_g\Delta} + \boldsymbol{M}_{s's,b>b'}^{U_g\Delta}|\boldsymbol{M}_{ss',b=b'}^{U_g\Delta} + \boldsymbol{M}_{s's,b=b'}^{U_g\Delta}\delta_{s<s'}\right], \ s \leq s' = 1, \ldots, \rho^{U_g}.$$

The constraint matrix $\boldsymbol{W}^{U_g}$ consists of $\left(\rho^{U_g}\right)^2 \times \frac{\left(\rho^{U_g}\right)^2+1}{2}$ blocks, most of which are zero. The block-rows of $\boldsymbol{W}^{U_g}$ correspond to the constraint on the spline coefficients of the covariances $K_{ss'}^{U_g}(t, t')$, $s, s' = 1, \ldots, \rho^{U_g}$, sorted as in $\boldsymbol{M}^{U_g\Delta}$. The columns are sorted correspondingly to the reduced matrix $\boldsymbol{M}^{U_g\Delta r}$. For the auto-covariances ($s = s'$), the blocks are of the same form as the constraint matrix $\boldsymbol{W}^E$ for independent curves. For the cross-covariances ($s < s'$), the blocks either correspond to diagonal block matrices or to anti-diagonal block matrices, depending on whether the respective rows correspond to $s < s'$ or $s > s'$, respectively. The specific form of $\boldsymbol{W}^{U_g}$ and examples for $\rho^{U_g} = 2, 3$ are given in Appendix B.

A quadratic smoothness penalty associated with each smooth term controls the bias-variance tradeoff. Each penalty matrix $\boldsymbol{S}^{U_g}$, consisting of blocks for each $K_{ss'}^{U_g}(t,t')$, is accordingly reduced by left-and right-multiplication with the constraint matrix $\boldsymbol{W}^{U_g}$. Smoothing the components in the upper triangle of $\boldsymbol{K}^{U_g}(t,t')$ separately, allows to define different penalties for the auto-covariances and the cross-covariances, respectively. In particular, it is possible to apply anisotropic penalties for the cross-covariances.

Reflecting the estimated triangular covariance surfaces across the diagonal yields estimates for the whole covariance surfaces $\boldsymbol{K}^{U_g}(t,t')$ and $K^E(t,t')$, with smoothness assured also across the diagonal. Note that smoothing multiple covariances using our approach reduces computation time compared to estimating all spline coefficients even more than in the independent case of Section 2.

3.3 Functional linear mixed model with crossed random intercepts

Motivated by our application to the phonetics data in Section 6.2, we now illustrate the specification of the FLMM for the special case of an FLMM with crossed fRIs accounting for the repeated measurements on two grouping variables (e.g. speakers and target words) in a crossed design. In this model, we have $G = 2$ grouping variables with $\rho^{U_1} = \rho^{U_2} = 1$ associated random effects for each grouping variable, i.e. one fRI each. L^{U_1}, L^{U_2} are the numbers of levels of the first (e.g. speakers) and second (e.g. target words) grouping variable, respectively. The covariate vector $\boldsymbol{z}_i$ only consists of indicators taking value one or zero to code group membership for the two grouping variables. The explicit specification of the covariate vector for crossed and hierarchical functional random effects is given in Appendix B. For better readability, we rename in the following the components of the vector-valued random process as $B := U_1$ and $C := U_2$. The total number of components in $\boldsymbol{U}(t_{ij})$ is $q = L^B + L^C$. Note that this model corresponds to the model in Cederbaum et al. (2016).

For this model, equation (3.3) can be simplified as

$$\begin{aligned} \mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\right] &= K^B(t_{ij},t_{i'j'})\delta_{B(i)B(i')} + K^C(t_{ij},t_{i'j'})\delta_{C(i)C(i')} \\ &+ \left[K^E(t_{ij},t_{i'j'}) + \sigma^2\delta_{jj'}\right]\delta_{ii'}, \end{aligned} \tag{3.6}$$

where $\delta_{B(i)B(i')}$ and $\delta_{C(i)C(i')}$ take value one when the two curves i and i' belong to the same level of the respective grouping variable and zero otherwise. As can be seen in Equation (3.6), the products for which neither $\delta_{B(i)B(i')}$ nor $\delta_{C(i)C(i')}$ equals one do not have to be considered due to expectation zero. Equation (3.4) then reduces to

$$\mathbb{E}\left(\boldsymbol{C}^{\Delta}\right) \;=\; \left[\boldsymbol{M}^{B\Delta}|\boldsymbol{M}^{C\Delta}|\boldsymbol{M}^{E\Delta}|\boldsymbol{\delta}^{\varepsilon\Delta}\right]\left(\boldsymbol{\theta}^{B\top},\boldsymbol{\theta}^{C\top},\boldsymbol{\theta}^{E\top},\sigma^2\right)^{\top} \tag{3.7}$$

and the symmetry constraint can be applied to each $\boldsymbol{\theta}^B$, $\boldsymbol{\theta}^C$, and $\boldsymbol{\theta}^E$.

3.4 Covariance of cross products

We estimate the auto-covariances as unknown, smooth functions in a bivariate additive varying coefficient model using a quadratic loss function. Since this is equivalent to a penalized likelihood criterion for Gaussian data, we implicitly assume independence of the cross products with homoscedastic Gaussian measurement error. As already mentioned in Section 2.2 and shortly discussed in Cederbaum et al. (2016), these are working assumptions which do not hold as the products frequently involve two points on the same curve or on correlated curves. Nevertheless, these implicit assumptions are made by many existing works (e.g. Yao et al., 2005; Di et al., 2009; Greven et al., 2010). We now briefly discuss how the covariance of the cross products can be accounted for.

For the model with independent curves (2.1), Xiao et al. (2016b) derive an expression for the covariance of the cross products in terms of $K^E(t,t')$ and the error variance σ^2 under the assumption of Gaussian responses. They apply a three-step algorithm in which they first estimate $K^E(t,t')$ and σ^2 under the working assumptions. Second, they estimate the covariance of the cross products by plugging in the estimates for $K^E(t,t')$ and σ^2. In the third step, they re-estimate $K^E(t,t')$ and σ^2 using the estimated covariance of the cross products as a working covariance.

We derive an expression for the covariance of the cross products for the general model (3.1) based on results from Isserlis (1918) on fourth moment rules for multivariate Gaussian

random variables. The covariance of the cross products can be written as

$$
\begin{aligned}
& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_m(t_{mo})\tilde{Y}_{m'}(t_{m'o'})\right] \qquad (3.8)\\
= \; & \left\{\boldsymbol{z}_i^\top \boldsymbol{K}^U(t_{ij}, t_{mo})\boldsymbol{z}_m + \left[K^E(t_{ij}, t_{mo}) + \sigma^2\delta_{jo}\right]\delta_{im}\right\} \\
& \cdot \left\{\boldsymbol{z}_{i'}^\top \boldsymbol{K}^U(t_{i'j'}, t_{m'o'})\boldsymbol{z}_{m'} + \left[K^E(t_{i'j'}, t_{m'o'}) + \sigma^2\delta_{j'o'}\right]\delta_{i'm'}\right\} \\
+ \; & \left\{\boldsymbol{z}_i^\top \boldsymbol{K}^U(t_{ij}, t_{m'o'})\boldsymbol{z}_{m'} + \left[K^E(t_{ij}, t_{m'o'}) + \sigma^2\delta_{jo'}\right]\delta_{im'}\right\} \\
& \cdot \left\{\boldsymbol{z}_{i'}^\top \boldsymbol{K}^U(t_{i'j'}, t_{mo})\boldsymbol{z}_m + \left[K^E(t_{i'j'}, t_{mo}) + \sigma^2\delta_{j'o}\right]\delta_{i'm}\right\}.
\end{aligned}
$$

The derivation and simplifications for the model with crossed fRIs are given in Appendix A. The covariance (3.8) is a function of the unknown covariances $\boldsymbol{K}^U(t, t')$, $K^E(t, t')$ and σ^2, giving rise to the need for an iterative procedure as proposed in Xiao et al. (2016b) for the simpler model (2.1). Our implementation with R-package mgcv allows to directly include the variance of the cross products by a specification of the weights argument in function bam. Our simulations showed, however, that accounting for only the heterogeneous variance does not lead to a substantial improvement in estimation accuracy. The dependencies could also be accounted for by pre-multiplication with the inverse square root of the covariance of the cross products, for which its construction and inversion would become necessary. For our application to the phonetics data, the covariance matrix of the cross products would be a $52,346,570 \times 52,346,570$ dense and unstructured matrix whose construction is not feasible with current technology ($\approx$ 17,500 Terabytes storage space would be required) and we thus do not focus on this extension in the following. However, (3.8) allows the inclusion of the covariance of the cross products in less complex settings and for (much) smaller data sets and can then result in more efficient estimates. Despite the violations of working assumptions, we achieve good results with our approach in simulations (cp. Section 7). Note that a relevant improvement of the covariance estimation can only be obtained if the working covariance is reasonably well estimated and thus more than one iteration of the algorithm might be necessary. Especially for multiple iterations, our fast smoothing algorithm considerably speeds up the estimation compared to smoothing the entire covariances.

4 Application in functional principal component analysis

An important application of our fast symmetric additive smoothing approach is functional principal component analysis. FPCA is a key tool for dimension reduction in FDA that extracts the dominant modes of variation in the data and provides an explicit variance decomposition. In the following, we briefly outline the four main steps of FPCA for the general FLMM (3.1) using our newly proposed covariance estimation approach. Applying FPCA to model (3.1) yields parsimonious representations of each random process $\boldsymbol{U}_g(t)$, $g = 1, \ldots, G$, and $E_i(t)$, in bases of eigenfunctions of the respective, previously estimated smooth auto-covariances. In addition, we briefly describe how our covariance smoothing approach can be combined with the general framework of functional additive mixed models (FAMM; Scheipl et al., 2015) allowing for approximate statistical inference for the mean conditional on the FPCA. For a more detailed description, see Cederbaum et al. (2016).

In the first step, the smooth mean function is commonly estimated based on a working independence assumption. We use penalized splines implemented in R-package mgcv, which can be used to estimate a large variety of covariate and interaction effects in the mean function. For the subsequent steps, the curves are then centered by subtracting the estimated mean from the functional observations.

For the second step, we propose to simultaneously estimate the upper triangles of the covariances $\boldsymbol{K}^{U_g}(t, t')$, $g = 1, \ldots, G$, and $K^E(t, t')$ and the error variance using our novel covariance smoothing approach. The triangular covariance surfaces are then reflected across the diagonal, yielding estimates for the complete covariance surfaces $\boldsymbol{K}^{U_g}(t, t') = \left[K^{U_g}_{ss'}(t, t')\right]_{s,s'=1,\ldots,\rho^{U_g}}$, $g = 1, \ldots, G$, and $K^E(t, t')$. Negative estimated values of σ^2 are set to zero.

In the third step, we use spectral decompositions of the estimated covariance surfaces based on Mercer's Theorem (Mercer, 1909)

$$\boldsymbol{K}^{U_g}(t, t') = \sum_{k=1}^{\infty} \nu_k^{U_g} \boldsymbol{\phi}_k^{U_g}(t) \boldsymbol{\phi}_k^{U_g}(t')^\top, \; \boldsymbol{K}^E(t, t') = \sum_{k=1}^{\infty} \nu_k^E \phi_k^E(t) \phi_k^E(t'),$$

with eigenvalues $\nu_k^{U_g}$, ν_k^E and (vector-valued) eigenfunctions $\boldsymbol{\phi}_k^{U_g}(t) = \left[\phi_{ks}^{U_g}(t)\right]_{s=1,\ldots,\rho^{U_g}}$, $\phi_k^E(t)$, respectively. In practice, the covariance surfaces are evaluated on a dense grid $T = \{t_1, \ldots, t_D\} \in \mathcal{T}$ of pre-specified length D. We obtain estimated eigenvalues $\hat{\nu}_k^{U_g}$, $k = 1, \ldots, D\rho^{U_g}$, and $\hat{\nu}_k^E$, $k = 1, \ldots, D$, as well as orthonormal eigenfunctions evaluated on T, $\hat{\boldsymbol{\phi}}_k^{U_g} = \left[\hat{\phi}_{ks}^{U_g}(t)\right]_{s=1,\ldots,\rho^{U_g},t\in T} \in \mathbb{R}^{D\rho^{U_g}}$ and $\hat{\boldsymbol{\phi}}_k^E = \left[\hat{\phi}_k^E(t)\right]_{t\in T} \in \mathbb{R}^D$, of each corresponding covariance operator. The multivariate eigenvectors $\hat{\boldsymbol{\phi}}_k^{U_g}$ consist of blocks for the respective eigenfunction components. The eigenvectors $\hat{\boldsymbol{\phi}}_k^{U_g}$, $\hat{\boldsymbol{\phi}}_k^E$ (and accordingly the eigenvalues) are rescaled to ensure orthonormality with respect to the additive scalar product $\langle (f_1, \ldots, f_{\rho^{U_g}}), (g_1, \ldots, g_{\rho^{U_g}}) \rangle = \sum_{s=1}^{\rho^{U_g}} \int_{\mathcal{T}} f_s(t) g_s(t) \,\mathrm{d}\, t$, and with respect to the L^2-scalar product $\langle f, g \rangle = \int_{\mathcal{T}} f(t) g(t) \,\mathrm{d}\, t$, respectively. For more details on multivariate FPCA, see e.g., Ramsay and Silverman (2005). To guarantee positive semi-definiteness of the covariances, which is not ensured by our smoothing approach, negative eigenvalues can be set to zero, which has been shown to improve accuracy in terms of the L^2-norm (Hall et al., 2008) and to work well in practice (Yao et al., 2003). Dimension reduction is achieved by truncating the number of eigenfunctions. We choose the truncation levels based on the proportion of variance explained (see Greven et al., 2010, for an overview) and denote them by N^{U_g}, $g = 1, \ldots, G$, and N^E, respectively. The truncated Karhunen-Loève (KL) expansion (Loève, 1945; Karhunen, 1947) allows parsimonious representations of the random processes in truncated bases of the corresponding eigenfunctions

$$\boldsymbol{U}_{gl}(t) \approx \sum_{k=1}^{N^{U_g}} \xi_{lk}^{U_g} \boldsymbol{\phi}_k^{U_g}(t), \;\; E_i(t) \approx \sum_{k=1}^{N^E} \xi_{ik}^E \phi_k^E(t), \tag{4.1}$$

with uncorrelated zero-mean random basis weights $\xi_{lk}^{U_g}$, $l = 1, \ldots, L^{U_g}$, $k = 1, \ldots, N^{U_g}$, and ξ_{ik}^E, $i = 1, \ldots, n$, $k = 1, \ldots, N^E$, with variance $\nu_k^{U_g}$ and ν_k^E, respectively.

In the fourth step, we predict the random basis weights, which give insight into the individual structure of each grouping level. Replacing the random processes in Model (3.1) by their truncated KL-expansions in (4.1) allows to approximate the model by a scalar linear mixed model with random effects corresponding to the random basis weights $\xi_{lk}^{U_g}$, ξ_{ik}^E (Di et al., 2009). The basis weights can then be predicted as empirical best linear unbiased predictors by simply plugging-in the estimated eigenfunctions, eigenvalues, and

the estimated error variance (Di et al., 2009; Greven et al., 2010; Cederbaum et al., 2016).

Alternatively, we can represent our model as a FAMM using our estimated eigenfunctions and -values in basis expansions of the random processes as proposed by Scheipl et al. (2015). The random basis weights are predicted together with a re-estimation of the mean function in a mixed model framework. This allows for more efficient mean estimation due to taking the covariance structure into account and for approximate statistical inference conditional on the FPCA, such as pointwise confidence bands for the mean and for covariate effects. For more details on the combination with the FAMM approach and an extensive comparison of the two ways to predict the basis weights, see Cederbaum et al. (2016).

5 Implementation

We base our implementation on R-package mgcv (Wood, 2011) which allows to add user-defined spline bases and penalties. In this framework, we define a novel class for bivariate smooths estimated subject to our symmetry constraint (2.5) called 'symm' by providing a new constructor method function smooth.construct.symm.smooth.spec and a corresponding predictor method function for the estimation of smooth surfaces in additive models. The class can be applied to any bivariate smooth term in a gam-formula. It is not restricted to symmetric data although we here consider symmetric data in the form of cross products in the smoothing of the covariance. It can be applied to (possibly noisy) data sampled on a regular grid as well as to irregularly or sparsely sampled data. As a modular component, our constructor can be applied separately to the auto-covariances of independent functional random effects in an FLMM. The case of correlated functional random effects is currently not covered in the implementation. In our application, we show how the constructor can be applied to complex designs on the basis of the FLMM with crossed fRIs as in (3.7). One main advantage of using standard software is that it allows for flexible extensions. The current implementation is based on tensor product B-splines with difference penalties (Eilers and Marx, 2003), but extensions to other bivariate bases that are symmetric across the diagonal and other penalties are possible.

The spline degree and the number of basis functions can be chosen. The user currently has the choice between two different quadratic penalties. One can either use the Kronecker

sum penalty $\text{pen}(\lambda) = \lambda\boldsymbol{\theta}^\top [(\boldsymbol{S}_t \otimes \boldsymbol{I}_F) + (\boldsymbol{I}_F \otimes \boldsymbol{S}_{t'})]\,\boldsymbol{\theta}$ or alternatively a Kronecker product penalty of the form $\text{pen}(\lambda) = \lambda\boldsymbol{\theta}^\top (\boldsymbol{S}_t \otimes \boldsymbol{S}_{t'})\,\boldsymbol{\theta} = \lambda\boldsymbol{\theta}^\top [(\boldsymbol{S}_t \otimes \boldsymbol{I}_F) \cdot (\boldsymbol{I}_F \otimes \boldsymbol{S}_{t'})]\,\boldsymbol{\theta}$, where $\boldsymbol{S}_t = \boldsymbol{S}_{t'}$ denote the marginal penalty matrices and $\boldsymbol{\theta}$ is the coefficient vector. Note that both penalties are isotropic, reflecting the symmetry of the surface. The main difference between the two penalty matrices is that the null space of the latter is larger, more likely leading to wigglier estimates. Other possible penalties could be added by the user.

To speed up estimation, function bam can compute the computationally expensive steps in parallel on multiple cores.

In addition to the R-code for our constructor, we provide code for the FPCA based on our (additive) covariance smoothing approach for three special cases of the FLMM (Model (2.1), a model with an fRI and a smooth error curve, and the model with two crossed fRIs and a smooth error curve as in Section 3.3).

6 Applications

We demonstrate the practical relevance of our approach in two distinct applications. We consider a standard data set consisting of sparse longitudinal observations as well as functional data with a complex correlation structure and different grids between curves.

6.1 CD4 cell count data

In order to compare our approach to that of Xiao et al. (2016b) for the common special case of longitudinal data, we apply FPCA to analyze the CD4 cell count trajectories in HIV positive individuals which are available in R-package refund (Huang et al., 2016). As our focus here is on the more complex case of correlated functional data, to which the approach of Xiao et al. (2016b) does not apply, the application to the CD4 cell count data is given in Appendix C.

6.2 Phonetics data

We apply FPCA based on our covariance smoother to acoustic signal data with a crossed correlation structure. We show the increase of the computational efficiency compared to the

approach of Cederbaum et al. (2016), for a case where their approach is applicable. To the best of our knowledge, their smooth method of moments approach is the only competitor for covariance smoothing of irregularly observed correlated curves with a crossed design structure. No alternative approach is available for general models (3.1), where our approach is the first available for additive covariance smoothing.

In phonetic research, the term consonant assimilation refers to the phenomenon that the articulation of two consonants becomes phonetically more alike when they appear subsequently in fluent speech. Consonant assimilation is accompanied by a complex interaction of language-specific, perceptual and articulatory factors which makes it an important topic in speech production research. The data we consider are part of a large study which was conducted by Pouplier and Hoole (2016) in order to investigate among others the assimilation of the consonants /s/ and /sh/ as a function of their order (/s#sh/ versus /sh#s/, where # denotes a word boundary), syllable stress and vowel context in the German language. The same data were previously analyzed by Cederbaum et al. (2016). Pouplier and Hoole (2016) recorded the audio signals for nine native speakers which repeated the same sixteen target words each five times. The target words consisted of (semantically nonsensical) bisyllabic noun-noun compound words with abutting consonants /s/ and /sh/ in either order, e.g. 'Calla**s-Sch**immel' and 'Gula**sch-S**impel', and with either stressed or unstressed syllables and varying vowel combinations. The time interval during the duration of the two consonants of interest was cut out manually by the phoneticians and standardized to a [0,1] interval in which the recorded acoustic signals were summarized in a functional index over time. The $n = 707$ index curves (shown in Figure 1) take values between $+1$ and -1, with positive [negative] values indicating proximity of the signal to a reference signal for the first [second] consonant of the target word, respectively. To illustrate the effect of consonant assimilation, two acoustic signals are highlighted in Figure 1. The curve without or with little assimilation shows a clear transition from a strong positive to a strong negative value, whereas the curve with strong assimilation is quite flat and mostly takes negative values. The curves differ in the number and location of the observation points, ranging from 22–57 with a median of 34 points per curve. For a more detailed description of the data (including pre-processing steps), see Pouplier and Hoole (2016); Cederbaum et al. (2016).

We fit an FLMM with crossed fRIs as described in Section 3.3 and previously consid-

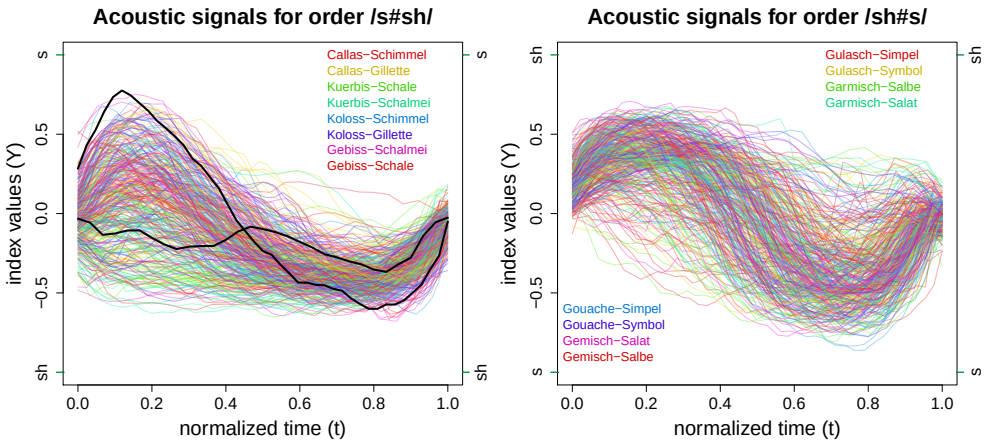

Figure 1: Acoustic signal curves of the phonetics data over time (cf. Cederbaum et al., 2016). Left [right]: Signal curves of consonant order /s#sh/ [sh#s] colored by target word. Example curves with no/little assimilation and with strong assimilation highlighted (black).

ered in Cederbaum et al. (2016) to account for the repeated measurements of both speakers and target words. The mean function $\mu(t, \boldsymbol{x}_i)$ includes effects and interaction effects of the dummy-coded covariates consonant order, syllable stress, and vowel context, smoothly varying over time. As the focus here is on the smooth auto-covariances of the functional random effects, $K^B(t, t')$, $K^C(t, t')$, and $K^E(t, t')$, we refer to Cederbaum et al. (2016) for more details on and interpretations of the covariate effects. For each auto-covariance, we use cubic marginal B-spline bases with marginal third order difference penalty matrices and use the Kronecker sum penalty for bivariate smoothing (cp. Section 5). We compare the results from our novel symmetric bivariate smoother (denoted by TRI-CONSTR and by TRI-CONSTR-W with weights of 0.5 on the cross products on the diagonal, cp. Section 2.2) with the results we obtain using the smoothing approach of Cederbaum et al. (2016) (denoted by WHOLE). The latter does not exploit the symmetry of the estimated surface and is equivalent to TRI-CONSTR-W except for the estimation of the smoothing parameter and numerical differences. To highlight the need for a symmetry constraint when only the upper triangle is considered, we further compare with the results obtained by estimating the upper triangle without a symmetry constraint (denoted by TRI) which does not guarantee

smoothness across the diagonal (cp. Section 2.2). The estimated auto-covariances are evaluated on a pre-specified grid of length $D = 100$. During the FPCA, we choose the truncation levels based on a pre-specified proportion of explained variance of 0.95, yielding two [four] eigenfunctions for the auto-covariance of the fRI for speakers, $K^B(t, t')$, and three [twelve] eigenfunctions for the auto-covariance of the smooth error, $K^E(t, t')$, for our two approaches and WHOLE [for TRI]. For all four smoothing methods, no eigenfunction is chosen for the fRI for target words, which is due to the high number of covariates that describe the target words sufficiently (cf. Cederbaum et al., 2016). Figure 2 shows the estimated surfaces and contours of the auto-covariance of the fRI for speakers, reconstructed after truncation from the estimated eigenvalues and eigenfunctions, which are shown in the bottom of the figure.

We can see from Figure 2 that the estimated covariance $\hat{K}^B(t, t')$ based on our symmetric smoother (TRI-CONSTR-W) is very similar to the one obtained by using all cross products (WHOLE). Moreover, it shows that for this application, the weights on the diagonal cross products do not make a great difference. As expected, we observe wigglier estimates especially on the diagonal for TRI, for which the error variance is estimated to be zero. This is also reflected in the wiggliness and higher number of chosen eigenfunctions for TRI. Similar results can be found for $K^E(t, t')$. These are given in Appendix C, where additional estimation details and results including the complete variance decompositions are provided. For all four approaches, the first eigenfunction in Figure 2 (solid line) corresponds to the discrimination of the speaker between the first and the second consonant and the second eigenfunction (dashed line) mainly leads to a vertical shift of the signal curves. Accounting for the symmetry of the covariances leads to a considerable reduction of computation times. TRI-CONSTR and TRI-CONSTR-W have the shortest computation times for smoothing the three auto-covariances, using five kernels in parallel, amounting to 24.51 and 25.82 minutes, respectively. Smoothing the auto-covariances using WHOLE takes more than twice as long (55.76 minutes) and using TRI still amounts to 32.95 minutes, which partly results from the fact that more spline coefficients have to be estimated. In addition, WHOLE and TRI require the estimation of two (instead of one) smoothing parameters for each auto-covariance using the Kronecker sum penalty implemented in R-package mgcv of the form $\text{pen}(\lambda) = \lambda_t \boldsymbol{\theta}^\top (\boldsymbol{S}_t \otimes \boldsymbol{I}_F) \boldsymbol{\theta} + \boldsymbol{\theta}^\top \lambda_{t'} (\boldsymbol{I}_F \otimes \boldsymbol{S}_{t'}) \boldsymbol{\theta}$.

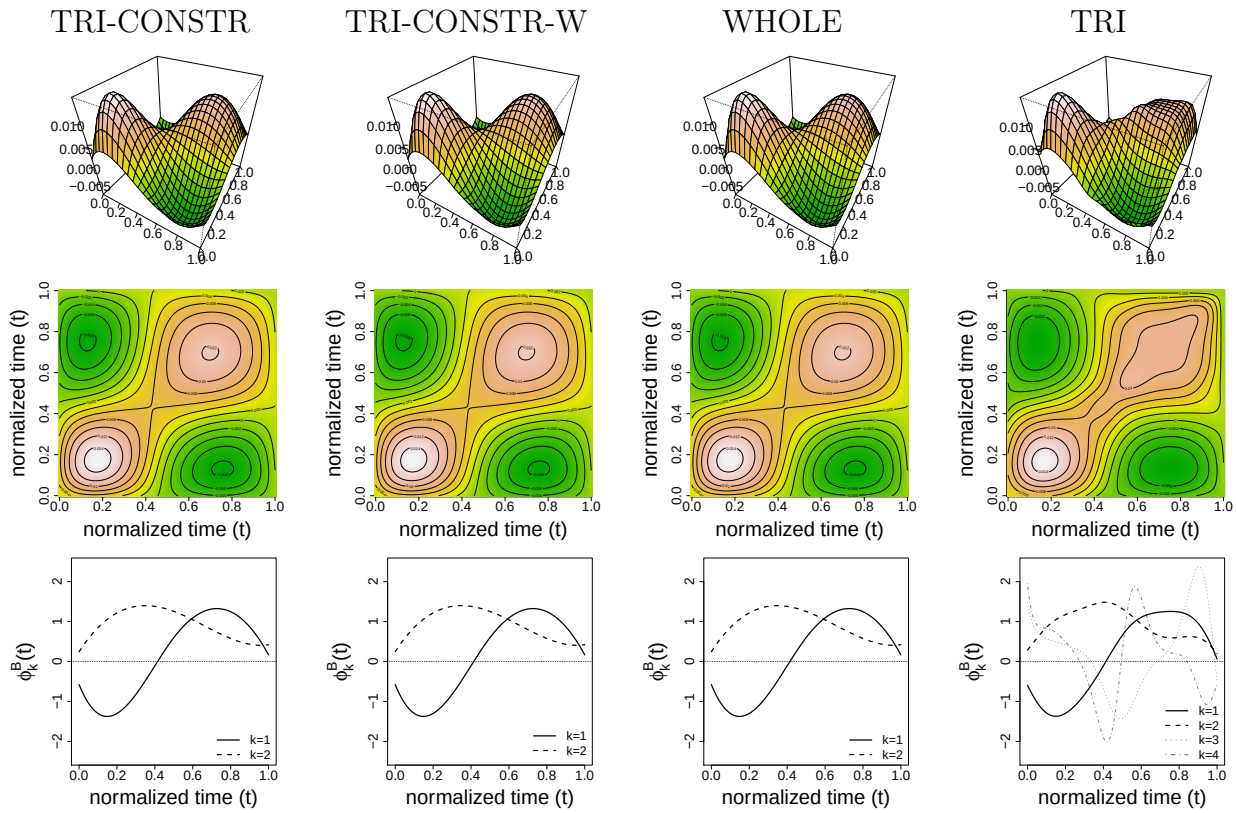

Figure 2: Results for the fRI for speakers using the four smoothing methods. Top row: estimated covariance surfaces. Middle row: contours of the estimated covariance surfaces. Bottom row: estimated corresponding eigenfunctions $\phi_k^B(t)$.

7 Simulations

7.1 Simulation designs

To investigate the performance of our covariance smoothing approach, we conduct an extensive simulation study based on two different data generating processes. For the first scenario (Scenario 1), we simulate data consisting of $n = 100$ independent curves from Model (2.1) with $\mu(t) = \sin(t) + t$. The data for the second scenario (Scenario 2) are generated from an FLMM with crossed fRIs as in Section 3.3, such that they mimic the

irregularly observed phonetics data. Note that no covariate mean effects are included such that we additionally obtain one eigenfunction for the fRI for target words and thus really have crossed fRIs. Additional generation details for this scenario are given in Appendix D. We compare the performance of our fast (additive) symmetric covariance smoother (TRI-CONSTR, TRI-CONSTR-W with weights of 0.5 on the diagonal cross products) as basis for FPCA with the performance of the smoothing approach proposed by Cederbaum et al. (2016) (WHOLE). To evaluate the need of the symmetry constraint when only the upper triangle is estimated, we further compare with the results obtained without posing a symmetry constraint (TRI). For Scenario 1, we additionally provide a comparison with the smoothing approach of Xiao et al. (2016b) (denoted by FACE), implemented in function face.sparse in R-package face (Xiao et al., 2016a). FACE does, however, not apply to additive covariance smoothing needed for correlated curves. Note that Xiao et al. (2016b) compare their symmetric covariance smoother to a number of other approaches that are all restricted to independent curves. They show that their approach is superior or comparable in terms of median integrated squared errors and inter quartile ranges (IQR) to the approach implemented in function fpca.sc (Di et al., 2009) in R-package refund based on bivariate B-splines with a difference penalty and to a self-coded variant based on thin plate regression splines for covariance smoothing. Moreover, they demonstrate their supremacy over the geometric likelihood approach of Peng and Paul (2009) and the local polynomial approach proposed in Yao et al. (2003). We thus do not include these alternatives here.

Based on Scenario 1, we investigate the sensitivity of the estimates to varying model complexity in terms of the complexity of the underlying eigenfunctions, signal to noise levels as functions of error variances and eigenvalues, and degree of sparseness. We consider all possible combinations of

1. simple eigenfunctions: $\{\phi_1(t) = 1, \phi_2(t) = \sqrt{3}(2t-1)\}$, complex eigenfunctions $\{\phi_1(t) = \sin(2\pi t), \phi_2(t) = \cos(2\pi t)\}$
2. error variance: $\sigma^2 = 0.5$, $\sigma^2 = 0.05$
3. eigenvalues: $\{\nu_1 = 0.15, \nu_2 = 0.075\}$, $\{\nu_1 = 2, \nu_2 = 1\}$
4. number of observation points: drawn from uniform distributions $\mathcal{U}[40, 60]$.

For the simple eigenfunctions and error variance $\sigma^2 \in \{0.05, 0.5\}$, we additionally consider a sparse setting, in which the number of observation points is drawn from the uniform

distribution $\mathcal{U}[3, 10]$ and for the complex eigenfunctions and eigenvalues $\nu_1 = 2, \nu_2 = 1$, we additionally reduce the value of the error variance to $\sigma^2 = 0.01$. For all settings, we generate 200 data sets. The random basis weights are centered and decorrelated such that the weights empirically have zero mean and a correlation of zero (see Cederbaum et al., 2016, for a discussion).

For the estimation in Scenario 1, we use ten cubic B-splines each for the estimation of the mean function and as marginal bases for the auto-covariances. For the estimation in Scenario 2, we use eight and five cubic B-splines for the estimation of the mean function and as marginal bases for the auto-covariances, respectively. We use Kronecker sum penalties (cp. Section 5) of marginal third order difference penalties for bivariate smoothing. Estimation of the smoothing parameter is based on REML, except for FACE, which uses leave-one-curve-out cross validation. We use equidistant knots in function face.sparse instead of the default (quantile based knots) which would require an adapted penalty that is not implemented. The arguments that determine the smoothing parameter search in function face.sparse are left at their defaults. As function face.sparse does not allow to specify a fixed truncation level, we choose the number of eigenfunctions based on a pre-specified proportion of explained variance of 0.95. Note that we use the proportion of explained variance in the functional observations, whereas Xiao et al. (2016a) use that in $E_i(t)$. In order to be able to differentiate the error incurred by the truncation of the covariance surface to a few leading FPCs from the pure covariance surface estimation error, we pre-specify the correct truncation lags for Scenario 2, for which no comparison with FACE is possible anyway.

7.2 Simulation results

We present and discuss the results of both simulation scenarios. For Scenario 1, we focus our presentation of the results on the setting with complex eigenfunctions, an error variance of $\sigma^2 = 0.05$ and eigenvalues of size $\nu_1 = 2$, $\nu_2 = 1$ (denoted as Setting 1). As a measure of goodness of fit, we use root relative mean squared errors (rrMSEs) of the form $\sqrt{(\text{true} - \text{estimated})^2/\text{true}^2}$. The complete results for all settings and the specific forms of the rrMSE for all model components, are given in Appendix D.

Figure 3 depicts boxplots of the rrMSEs for 200 simulation runs for Setting 1. For

each model component, it shows the boxplots for the compared smoothing methods and, in addition, for a modified version of FACE (denoted by FACE-STEP-1), in which the covariance of the cross products is not accounted for and thus only the first step of the three-step procedure is performed. FACE-STEP-1 is added to evaluate the effect of accounting for the covariance of the cross products (cf. Section 3.4). It shows that all components, except for the error variance, are estimated very well for our approach (TRI-CONSTR, TRI-CONSTR-W). The weights on the diagonal cross products do not have a great influence. Moreover, our approach yields similar rrMSEs as WHOLE. We also obtain similar results for TRI for most components, except for the error variance for which TRI has a higher median rrMSE. FACE yields a more than 2.7 times higher median rrMSE for the auto-covariance $K^E(t, t')$ compared to our approach and consequently also worse results for the eigenvalues ν_k^E, $k = 1, 2$. It yields smaller median rrMSEs for the eigenfunctions ϕ_k^E, $k = 1, 2$, but the IQR is larger and more outliers occur than in all other approaches. The estimation of the error variance profits most from accounting for the covariance of the cross products, which is reflected in a much lower median rrMSE for FACE than for the other methods. This also leads to lower median rrMSEs for the random basis weights, ξ_1^E and ξ_2^E, which depend on the error variance and consequently also to lower median rrMSEs for the reconstructed processes $E_i(t)$ and $Y_i(t)$. For all components, however, FACE yields a high number of outliers that range above the maximal rrMSEs obtained for the other methods. It is noticeable that the estimation of the auto-covariance is considerably better (in terms of rrMSE) for FACE when the covariance of the cross products is not accounted for. Moreover, we see that FACE-STEP-1 yields roughly similar results to our approach and WHOLE. For TRI-CONSTR, TRI-CONSTR-W and WHOLE, the truncation level is correctly estimated to be two in all 200 simulation runs. A higher number of eigenfunctions (three to four) is chosen for TRI in eight simulation runs of this setting which corresponds to our results in the application to the phonetics data. FACE and FACE-STEP-1 choose more than two (three to five) eigenfunctions in 190 and 198 simulation runs, respectively.

To sum up the results for the other ten settings of Scenario 1, we can say that over all settings and components TRI-CONSTR, TRI-CONSTR-W and WHOLE yield pretty similar rrMSEs with a tendency to a supremacy of TRI-CONSTR, especially in the sparse settings. TRI yields similar to worse results compared to our method and WHOLE. Espe-

cially for the error variance it yields up to 82% higher median rrMSEs (in one of the sparse settings). The estimation quality of all methods differs between the dense and the sparse settings. Our approach yields relatively similar results within the dense settings and higher rrMSEs in the sparse settings. Moreover, TRI-CONSTR and TRI-CONSTR-W tend to perform better in the settings with complex eigenfunctions and favor smaller error variances (except for the estimation quality of the error variance itself). In contrast, FACE tends to perform better in the settings with simple eigenfunctions and favors larger error variances. Our approach and WHOLE always select the correct truncation level, except in the sparse settings, where for some simulation runs more eigenfunctions are selected. TRI tends to select more eigenfunctions. For all settings, FACE and FACE-STEP-1 have simulation runs in which more than two eigenfunctions are selected.

In contrast to FACE and FACE-STEP-1, our approach is directly extendable to smoothing multiple auto-covariances simultaneously, which is required in Scenario 2. Figure 4 depicts the rrMSEs for the auto-covariances, the error variance and curves $Y_i(t)$ for Scenario 2. It shows that the three auto-covariances are estimated equally well for the four compared smoothing methods. For the error variance, however, TRI performs worse with a 32.5% higher median rrMSE and a larger IQR which also results in slightly higher rrMSEs for the reconstructed curves. Over all, our methods perform very well considering the small number of levels for the fRIs B (9 levels) and C (16 levels). For all model components, TRI-CONSTR and TRI-CONSTR-W yield similar rrMSEs.

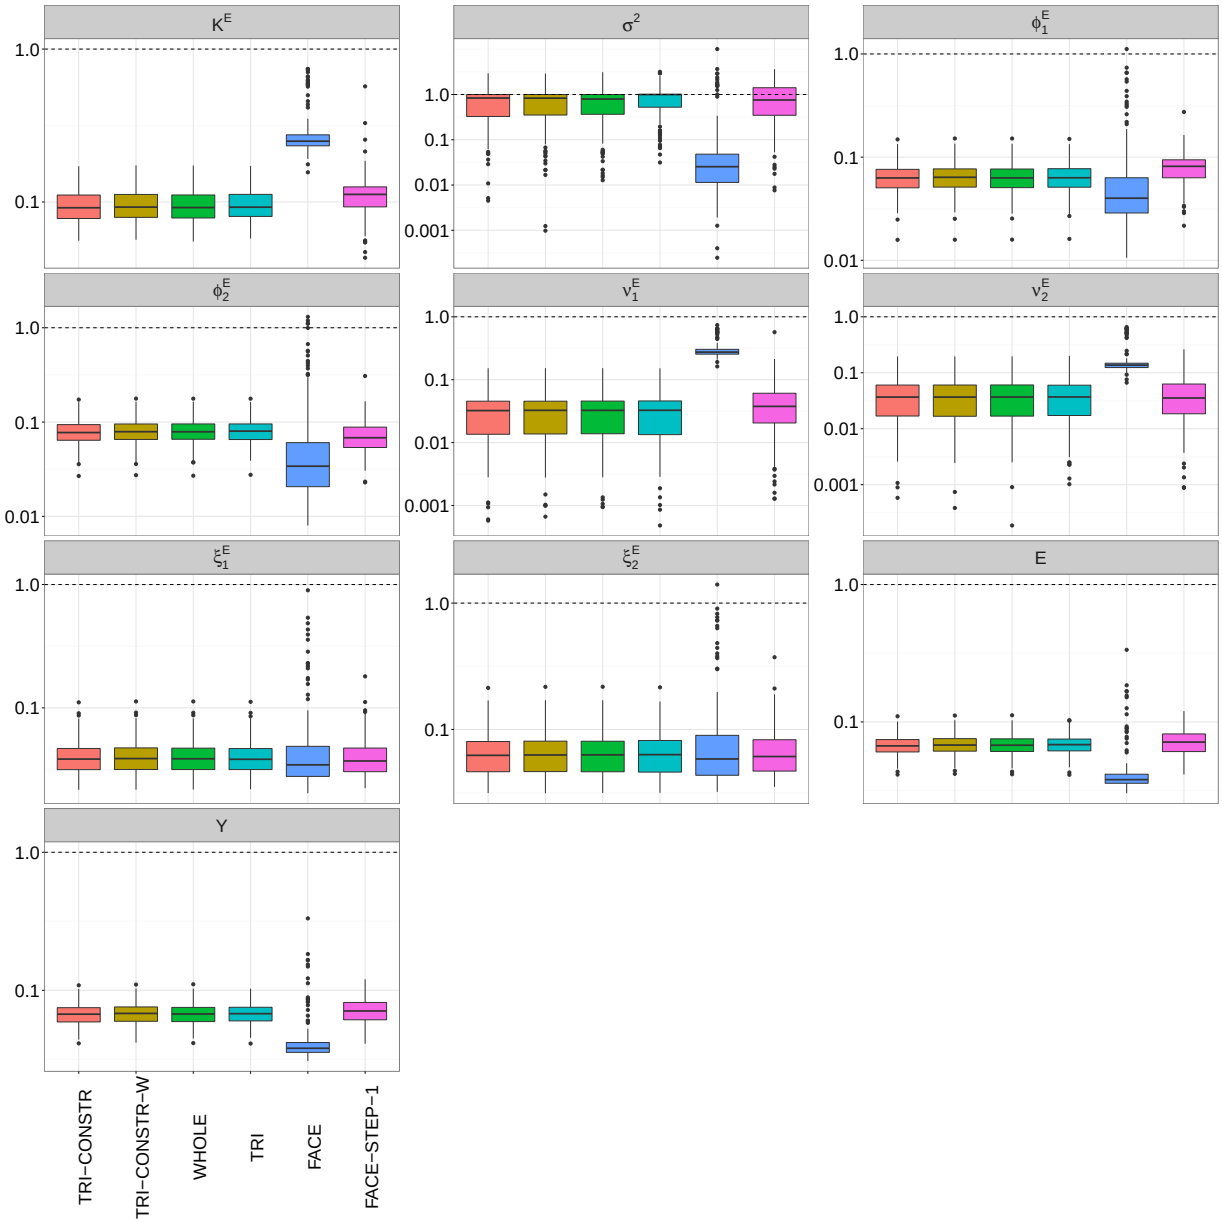

Figure 3: Boxplots of the rrMSEs (log10 scale at y-axis) for the scenario with independent curves (Scenario 1) for the smoothing methods being compared. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

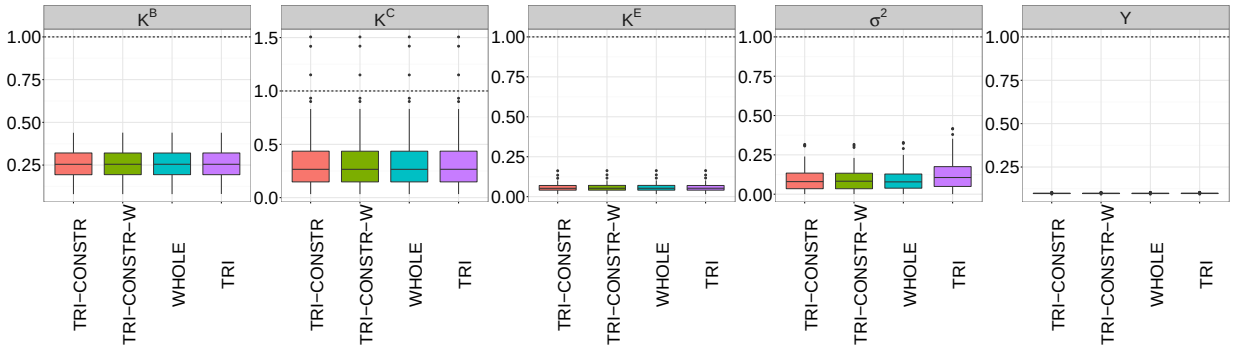

Figure 4: Boxplots of the rrMSEs for the scenario with crossed fRIs (Scenario 2) for the smoothing methods being compared. Shown are boxplots of the rrMSEs for the three auto-covariances $K^B(t,t')$, $K^C(t,t')$, and $K^E(t,t')$ as well as for the error variance σ^2 and the curves $Y_i(t)$.

7.3 Computational efficiency

Figure 5 shows computation times on a 64 Bit Linux platform with 660 Gb of RAM memory for the two settings discussed above. Our approach (with and without weights on the diagonal cross products) greatly speeds up the computation compared to WHOLE and also to TRI, especially in the scenario with crossed fRIs (right figure), where the computation times are longer and thus matter more. Note that in addition to accounting for the symmetry of covariances, we reduce the number of smoothing parameters to be estimated compared to WHOLE and TRI, for which two smoothing parameters are estimated for each auto-covariance using the Kronecker sum penalty implemented in R-package mgcv. In the scenario with independent curves (left figure), FACE is considerably slower than the other methods. This applies to all except the sparse settings, in which the computation times are extremely short anyway (max. 15 sec.). Note that variation for FACE and FACE-STEP-1 is considerably smaller as the smoothing parameter is chosen based on a fixed grid.

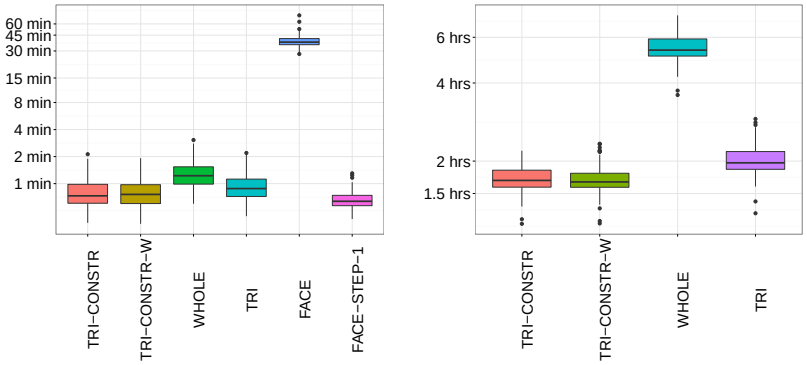

Figure 5: Computation times (log10 scale at y-axis) for 200 simulation runs for all compared methods. Left: computation times for Setting 1 of the scenario with independent curves. Right: computation times for the scenario with crossed fRIs.

8 Discussion and Outlook

We have introduced a fast bivariate smoothing approach for symmetric surfaces which applies to a broad range of data situations. We focus on its application to estimate covariance functions in longitudinal data as well as multiple additive covariance functions in functional data with very general correlation structures. Our smoother can handle (possibly noisy) data sampled on a common, dense grid as well as irregularly or sparsely observed data, which are frequently encountered in practice. It extends the smooth methods of moments estimator of Cederbaum et al. (2016) to more general correlation structures and additionally takes advantage of the symmetry of the sample covariances, which leads to considerably faster estimation requiring less memory. A symmetry constraint additionally ensures smoothness of the estimated covariance surfaces across the diagonal and further reduces computational costs. We show how our smoother can be applied as basis for FPCA, a key tool for dimension reduction in FDA, and demonstrate its practical relevance in a longitudinal data application and in an application to complexly correlated functional phonetics data.

We provide software implementing our approach that builds on the established R-package mgcv Wood (2011) allowing for flexible extensions. Within this framework, we provide a novel constructor function for the estimation of smooth surfaces in additive models subject

to the symmetry constraint. Our constructor can be applied as a modular component to general bivariate symmetric smoothing problems.

Simulation experiments (in Section 7) show that the proposed method recovers the true functions very well and yields similar results as the estimation approach of Cederbaum et al. (2016) while considerably speeding up the estimation and extending the range of possible model structures.

This work opens up a number of interesting directions for future research. A first direction concerns the working assumptions of the covariance estimation as additive varying coefficient model for the cross products (cp. Section 3.4). It would be interesting to investigate whether a suitable loss function for the cross products could be derived. Under the assumption of Gaussian responses, the cross products follow a product normal distribution, for whose PDF Nadarajah and Pogány (2016) recently derived a closed-form expression based on the modified Bessel function. A second direction concerns the positive semi-definiteness of the covariance operator, which is not ensured in our approach and in most existing covariance smoothing approaches (e.g. Yao et al., 2003; Hall et al., 2008; Di et al., 2009; Greven et al., 2010). Although Hall et al. (2008) show that setting negative eigenvalues to zero improves the estimation quality and Yao et al. (2003) demonstrate that this works well in practice, it could be desirable to ensure positive semi-definiteness in the estimation. Wu and Pourahmadi (2003) estimate a positive semi-definite covariance based on an auto-regressive model with regression coefficients corresponding to the components of a modified Cholesky decomposition of the covariance. Also the approach of Peng and Paul (2009) ensures positive semi-definiteness. It remains an open question, however, how these approaches could be extended to smoothing multiple additive covariances for functional data with complex correlation structures.

Acknowledgements

The authors thank Marianne Pouplier and Phil Hoole for supplying and explaining the phonetics data. All authors were supported by the Emmy Noether grant GR 3793/1-1 from the German Research Foundation.

References

Besse, P. and Ramsay, J. O. (1986). Principal components analysis of sampled functions. *Psychometrika*, 51(2):285–311.

Cederbaum, J., Pouplier, M., Hoole, P., and Greven, S. (2016). Functional linear mixed models for irregularly or sparsely sampled data. *Statistical Modelling*, 16(1):67–88.

Di, C.-Z., Crainiceanu, C. M., Caffo, B. S., and Punjabi, N. M. (2009). Multilevel functional principal component analysis. *The Annals of Applied Statistics*, 3(1):458–488.

Di, C.-Z., Crainiceanu, C. M., and Jank, W. S. (2014). Multilevel sparse functional principal component analysis. *Stat*, 3(1):126–143.

Diggle, P., Heagerty, P., Liang, K.-Y., and Zeger, S. (2002). *Analysis of longitudinal data*. Oxford University Press, Oxford, United Kingdom.

Eilers, P. H. and Marx, B. D. (2003). Multivariate calibration with temperature interaction using two-dimensional penalized signal regression. *Chemometrics and Intelligent Laboratory Systems*, 66(2):159–174.

Fan, J., Huang, T., and Li, R. (2007). Analysis of longitudinal data with semiparametric estimation of covariance function. *Journal of the American Statistical Association*, 102(478):632–641.

Faraway, J. J. (1997). Regression analysis for a functional response. *Technometrics*, 39(3):254–261.

Goldsmith, J., Greven, S., and Crainiceanu, C. M. (2013). Corrected confidence bands for functional data using principal components. *Biometrics*, 69(1):41–51.

Greven, S., Crainiceanu, C. M., Caffo, B., and Reich, D. (2010). Longitudinal functional principal component analysis. *Electronic Journal of Statistics*, 4:1022–1054.

Hall, P., Müller, H., and Yao, F. (2008). Modelling sparse generalized longitudinal observations with latent Gaussian processes. *Journal of the Royal Statistical Society: Series B*, 70(4):703–723.

Huang, L., Scheipl, F., Goldsmith, J., Gellar, J., Harezlak, J., McLean, M. W., Swihart, B., Xiao, L., Crainiceanu, C., and Reiss, P. (2016). *refund: Regression with Functional Data*. R package version 0.1-14. Available from: `https://CRAN.R-project.org/package=refund`.

Isserlis, L. (1918). On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. *Biometrika*, 12(1/2):134–139.

James, G. M. and Hastie, T. J. (2001). Functional linear discriminant analysis for irregularly sampled curves. *Journal of the Royal Statistical Society: Series B*, 63(3):533–550.

James, G. M., Hastie, T. J., and Sugar, C. A. (2000). Principal component models for sparse functional data. *Biometrika*, 87(3):587–602.

Karhunen, K. (1947). Über lineare Methoden in der Wahrscheinlichkeitsrechnung. *Annales Academiae Scientiarum Fennicae*, 37:1–79.

Kaslow, R. A., Ostrow, D. G., Detels, R., Phair, J. P., Polk, B. F., Rinaldo, C. R., et al. (1987). The multicenter aids cohort study: rationale, organization, and selected characteristics of the participants. *American journal of epidemiology*, 126(2):310–318.

Kauermann, G. and Wegener, M. (2011). Functional variance estimation using penalized splines with principal component analysis. *Statistics and Computing*, 21(2):159–171.

Krivobokova, T. and Kauermann, G. (2007). A note on penalized spline smoothing with correlated errors. *Journal of the American Statistical Association*, 102(480):1328–1337.

Leurgans, S. E., Moyeed, R. A., and Silverman, B. W. (1993). Canonical correlation analysis when the data are curves. *Journal of the Royal Statistical Society. Series B*, 55:725–740.

Loève, M. (1945). Fonctions aléatoires du second ordre. *Comptes Rendus de l'Académie des Sciences. Série I: Mathématique*, 220:469.

Mercer, J. (1909). Functions of positive and negative type, and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society of London. Series A*, 209:415–446.

Morris, J. S. (2015). Functional regression. *Annual Review of Statistics and Its Application*, 2(1):321–359.

Nadarajah, S. and Pogány, T. K. (2016). On the distribution of the product of correlated normal random variables. *Comptes Rendus Mathematique*, 354(2):201–204.

Patterson, H. D. and Thompson, R. (1971). Recovery of inter-block information when block sizes are unequal. *Biometrika*, 58(3):545–554.

Peng, J. and Paul, D. (2009). A geometric approach to maximum likelihood estimation of the functional principal components from sparse longitudinal data. *Journal of Computational and Graphical Statistics*, 18(4):995–1015.

Pinheiro, J. C. and Bates, D. M. (2000). *Mixed-Effects Models in S and S-PLUS*. Springer, New York.

Pouplier, M. and Hoole, P. (2016). Articulatory and acoustic characteristics of german fricative clusters. *Phonetica*, 73(1):52–78.

R Development Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer, New York, 2nd edition.

Reiss, P. T., Huang, L., and Mennes, M. (2010). Fast function-on-scalar regression with penalized basis expansions. *The International Journal of Biostatistics*, 6(1):28.

Scheipl, F., Staicu, A.-M., and Greven, S. (2015). Functional additive mixed models. *Journal of Computational and Graphical Statistics*, 24(2):477–501.

Shou, H., Zipunnikov, V., Crainiceanu, C. M., and Greven, S. (2015). Structured functional principal component analysis. *Biometrics*, 71(1):247–257.

Staniswalis, J. and Lee, J. (1998). Nonparametric regression analysis of longitudinal data. *Journal of the American Statistical Association*, 93(444):1403–1404.

Tai, W. (2009). *Regularized estimation of covariance matrices for longitudinal data through smoothing and shrinkage*. PhD thesis, Columbia University.

Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B*, 73(1):3–36.

Wu, W. B. and Pourahmadi, M. (2003). Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika*, 90(4):831–844.

Xiao, L., Li, C., Checkley, W., and Crainiceanu, C. (2016a). *face: Fast Covariance Estimation for Sparse Functional Data*. R package version 0.1-1. Available from: `https://CRAN.R-project.org/package=face`.

Xiao, L., Li, C., Checkley, W., and Crainiceanu, C. M. (2016b). Fast covariance estimation for sparse functional data. *ArXiv e-prints: 1603.05758v1*. Available from: `http://adsabs.harvard.edu/abs/2016arXiv160305758X`.

Xiao, L., Zipunnikov, V., Ruppert, D., and Crainiceanu, C. M. (2016c). Fast covariance estimation for high-dimensional functional data. *Statistics and Computing*, 26(1):409–421.

Yao, F., Müller, H.-G., Clifford, A. J., Dueker, S. R., Follett, J., Lin, Y., Buchholz, B. A., and Vogel, J. S. (2003). Shrinkage estimation for functional principal component scores with application to the population kinetics of plasma folate. *Biometrics*, 59(3):676–685.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590.

A Derivations

A.1 Derivation for the covariance of cross products in the general FLMM

Let $X = (X_1, \ldots, X_p)^\top$ be a p-dimensional Gaussian random variable with zero mean and covariance $\boldsymbol{\Sigma}$. Then, given $p \geq 4$, we can express the fourth moment of X based on Isserlis' Theorem (Isserlis, 1918) as

$$\mathbb{E}\left(X_i X_j X_k X_l\right) = \boldsymbol{\Sigma}_{ij}\boldsymbol{\Sigma}_{kl} + \boldsymbol{\Sigma}_{ik}\boldsymbol{\Sigma}_{jl} + \boldsymbol{\Sigma}_{il}\boldsymbol{\Sigma}_{jk}, \tag{A.1}$$

where $\boldsymbol{\Sigma}_{ij}$ is the covariance of X_i and X_j.

Consider the covariance of the cross products of the centered functional responses in the general FLMM

$$\begin{aligned}
& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_m(t_{mo})\tilde{Y}_{m'}(t_{m'o'})\right] \qquad (A.2)\\
= \;& \mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\tilde{Y}_m(t_{mo})\tilde{Y}_{m'}(t_{m'o'})\right] - \underbrace{\mathbb{E}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\right]}_{\operatorname{Cov}\left[\tilde{Y}_i(t_{ij}),\tilde{Y}_{i'}(t_{i'j'})\right]}\underbrace{\mathbb{E}\left[\tilde{Y}_m(t_{mo})\tilde{Y}_{m'}(t_{m'o'})\right]}_{\operatorname{Cov}\left[\tilde{Y}_m(t_{mo}),\tilde{Y}_{m'}(t_{m'o'})\right]}\\
\overset{(A.1)}{=} \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'})\right]\operatorname{Cov}\left[\tilde{Y}_m(t_{mo}), \tilde{Y}_{m'}(t_{m'o'})\right]\\
+ \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_m(t_{mo})\right]\operatorname{Cov}\left[\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_{m'}(t_{m'o'})\right]\\
+ \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{m'}(t_{m'o'})\right]\operatorname{Cov}\left[\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_m(t_{mo})\right]\\
- \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij}), \tilde{Y}_{i'}(t_{i'j'})\right]\operatorname{Cov}\left[\tilde{Y}_m(t_{mo}), \tilde{Y}_{m'}(t_{m'o'})\right]\\
= \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_m(t_{mo})\right]\operatorname{Cov}\left[\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_{m'}(t_{m'o'})\right]\\
+ \;& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{m'}(t_{m'o'})\right]\operatorname{Cov}\left[\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_m(t_{mo})\right]\\
= \;& \left\{\boldsymbol{z}_i^\top \boldsymbol{K}^U(t_{ij}, t_{mo})\boldsymbol{z}_m + \left[K^E(t_{ij}, t_{mo}) + \sigma^2\delta_{jo}\right]\delta_{im}\right\}\\
& \cdot \left\{\boldsymbol{z}_{i'}^\top \boldsymbol{K}^U(t_{i'j'}, t_{m'o'})\boldsymbol{z}_{m'} + \left[K^E(t_{i'j'}, t_{m'o'}) + \sigma^2\delta_{j'o'}\right]\delta_{i'm'}\right\}\\
+ \;& \left\{\boldsymbol{z}_i^\top \boldsymbol{K}^U(t_{ij}, t_{m'o'})\boldsymbol{z}_{m'} + \left[K^E(t_{ij}, t_{m'o'}) + \sigma^2\delta_{jo'}\right]\delta_{im'}\right\}\\
& \cdot \left\{\boldsymbol{z}_{i'}^\top \boldsymbol{K}^U(t_{i'j'}, t_{mo})\boldsymbol{z}_m + \left[K^E(t_{i'j'}, t_{mo}) + \sigma^2\delta_{j'o}\right]\delta_{i'm}\right\}. \quad \square
\end{aligned}$$

A.2 Simplification of the covariance of cross products for crossed fRIs

For the special case of an FLMM with two crossed fRIs as in Section 3.3, the covariance of the cross products in equation (A.2) simplifies to

$$
\begin{aligned}
& \operatorname{Cov}\left[\tilde{Y}_i(t_{ij})\tilde{Y}_{i'}(t_{i'j'}), \tilde{Y}_m(t_{mo})\tilde{Y}_{m'}(t_{m'o'})\right] \\
= \; & \{K^B(t_{ij}, t_{mo})\delta_{B(i)B(m)} + K^C(t_{ij}, t_{mo})\delta_{C(i)C(m)} + \left[K^E(t_{ij}, t_{mo}) + \sigma^2\delta_{jo}\right]\delta_{im}\} \\
\cdot \; & \{K^B(t_{i'j'}, t_{m'o'})\delta_{B(i')B(m')} + K^C(t_{i'j'}, t_{m'o'})\delta_{C(i')C(m')} + \left[K^E(t_{i'j'}, t_{m'o'}) + \sigma^2\delta_{j'o'}\right]\delta_{i'm'}\} \\
+ \; & \{K^B(t_{ij}, t_{m'o'})\delta_{B(i)B(m')} + K^C(t_{ij}, t_{m'o'})\delta_{C(i)C(m')} + \left[K^E(t_{ij}, t_{m'o'}) + \sigma^2\delta_{jo'}\right]\delta_{im'}\} \\
\cdot \; & \{K^B(t_{i'j'}, t_{mo})\delta_{B(i')B(m)} + K^C(t_{i'j'}, t_{mo})\delta_{C(i')C(m)} + \left[K^E(t_{i'j'}, t_{mo}) + \sigma^2\delta_{j'o}\right]\delta_{i'm}\},
\end{aligned}
$$

where $\delta_{B(i)B(m)}$ and $\delta_{C(i)C(m)}$ take value one when the two curves i and m belong to the same level of the respective grouping variable and zero otherwise.

B Supplementary details on the estimation and implementation

B.1 (Additive) varying coefficient model using tensor product B-splines

Let $\otimes$ and $\cdot$ denote the Kronecker product and the Hadamard (pointwise) product, respectively.

Model with independent curves Using tensor product B-splines yields the following form of Model (2.3)

$$\begin{aligned}\mathbb{E}[\boldsymbol{C}] &= \left[\left(\boldsymbol{B}_t^E \otimes \mathbb{1}_{F^E}{}^\top\right) \cdot \left(\mathbb{1}_{F^E}{}^\top \otimes \boldsymbol{B}_{t'}^E\right) | \boldsymbol{\delta}^\varepsilon\right] \left(\boldsymbol{\theta}^{E^\top}, \sigma^2\right)^\top \\ &= \left[\boldsymbol{M}^E | \boldsymbol{\delta}^\varepsilon\right] \left(\boldsymbol{\theta}^{E^\top}, \sigma^2\right)^\top = \boldsymbol{M}\boldsymbol{\alpha},\end{aligned}$$

where $\boldsymbol{B}_t^E$, $\boldsymbol{B}_{t'}^E$ are the $\mathcal{C} \times F^E$ marginal spline design matrices that contain the evaluated spline basis functions for the directions t and t', respectively. $\boldsymbol{B}_t^E$ and $\boldsymbol{B}_{t'}^E$ contain identical but permuted rows. $\mathbb{1}_{F^E} = (1, \ldots, 1)^\top$ is of length F^E, the number of marginal basis functions in each direction. The bivariate additive model for the reduced response vector is obtained by replacing the marginal spline design matrices by the reduced $\mathcal{C}^\Delta \times F^E$ matrices $\boldsymbol{B}_t^{E\Delta}$, $\boldsymbol{B}_{t'}^{E\Delta}$ and the index vector by the reduced index vector $\boldsymbol{\delta}^{\varepsilon\Delta}$. Note that in the model with independent curves, the bivariate spline design matrix $\boldsymbol{B}^{E\Delta}$ corresponds to the design matrix $\boldsymbol{M}^{E\Delta}$ as products are only computed on the same curves and thus the indicator matrix $\boldsymbol{Q}^{E\Delta}$ reduces to an all-ones matrix.

General FLMM For each $g = 1, \ldots, G$, and each $s, s' = 1, \ldots \rho^{U_g}$, let $\boldsymbol{B}_{ss',t}^{U_g\Delta}$ and $\boldsymbol{B}_{ss',t'}^{U_g\Delta}$ denote the marginal spline design matrices of dimensions $\mathcal{C}^\Delta \times F_{ss',t}^{U_g}$ and $\mathcal{C}^\Delta \times F_{ss',t'}^{U_g}$, respectively. Due to the symmetry of the covariances $\boldsymbol{K}^{U_g}(t, t')$, we assume $F_{ss',t}^{U_g} = F_{s's,t'}^{U_g}$ and

$F^{U_g}_{ss',t'} = F^{U_g}_{s's,t}$, respectively. Then, the bivariate spline design matrices $\boldsymbol{B}^{U_g\Delta}_{ss'}$ are given by

$$\boldsymbol{B}^{U_g\Delta}_{ss'} \;=\; \left(\boldsymbol{B}^{U_g\Delta}_{ss',t} \otimes \mathbb{1}_{F^{U_g}_{ss',t'}}{}^{\top}\right) \cdot \left(\mathbb{1}_{F^{U_g}_{ss',t}}{}^{\top} \otimes \boldsymbol{B}^{U_g\Delta}_{ss',t'}\right).$$

The submatrices $\boldsymbol{M}^{U^\Delta_g}_{ss'}$ corresponding to the covariances $K^{U_g}_{ss'}(t,t')$ are given as $\boldsymbol{M}^{U_g\Delta}_{ss'} = \boldsymbol{Q}^{U^\Delta_g}_{ss'} \cdot \boldsymbol{B}^{U^\Delta_g}_{ss'}$, where $\boldsymbol{Q}^{U_g\Delta}_{ss'}$ are $\mathcal{C}^\Delta \times F^{U_g}_{ss',t}F^{U_g}_{ss',t'}$ matrices with entries $\delta_{U_g(i)U_g(i')} \cdot \omega^{U_g}_{is}\omega^{U_g}_{i's'}$, $i \leq i' = 1,\ldots,n$. $\delta_{U_g(i)U_g(i')}$ takes value one if the two curves i and i' are of the same level of grouping variable g and zero otherwise. The columns of $\boldsymbol{Q}^{U_g\Delta}_{ss'}$ are all identical and contain the suitably sorted and repeated entries. Suitably sorted and repeated in this context means that the sorting corresponds to the sorting in $\boldsymbol{C}^\Delta$ and that the entries $\delta_{U_g(i)U_g(i')} \cdot \omega^{U_g}_{is}\omega^{U_g}_{i's'}$ are repeated for all considered combinations of observation points $t_{ij} \leq t_{i'j'}$.

In analogy, matrix $\boldsymbol{Q}^{E\Delta}$ is a $\mathcal{C}^\Delta \times \left(F^E\right)^2$ matrix with identical columns consisting of suitably sorted and repeated indicators $\delta_{E(i)E(i')}$, which take value one if the two points in the cross products belong to the same curve and zero otherwise.

B.2 Examples for the general FLMM

In the following, we provide examples for the specification of $\boldsymbol{z}_i^\top \boldsymbol{U}(t_{ij})$ in Model (3.1) yielding an FLMM with hierarchical (e.g. subjects in groups) and crossed (e.g. speakers and target words as in our phonetics application in Section (6.2)) functional random effects, respectively.

Consider for simplicity the case of two independent grouping variables ($G = 2$), with an fRI for the first grouping variable and an fRI and a functional random slope (in variable ω) for the second grouping variable ($\rho^{U_1} = 1$, $\rho^{U_1} = 2$). Further assume that there are n_i observations (curves) for each level of the second grouping variable, $i = 1,\ldots,L^{U_2}$. We assume in the following that the fRIs and the functional random slopes are dependent. For independent effects, we would split them up and denote $G = 3$. Let ω_i denote the value of variable ω for curve i and $\delta_{U_g(i)l}$ takes value one if curve i belongs to the lth level of

grouping variable g and zero otherwise. Then, $\boldsymbol{z}_i^\top \boldsymbol{U}(t_{ij})$, $i = 1, \ldots, n$, is given by

$$\boldsymbol{z}_i^\top \boldsymbol{U}(t_{ij}) = \left(\delta_{i11}^U, \ldots, \delta_{i1L^{U_1}}^U, \delta_{i21}^U, \delta_{i21}^U \omega_i, \ldots, \delta_{i2L^{U_2}}^U, \delta_{i2L^{U_2}}^U \omega_i\right) \begin{pmatrix} U_{111}(t_{ij}) \\ \vdots \\ U_{1L^{U_1}1}(t_{ij}) \\ \vdots \\ U_{211}(t_{ij}) \\ U_{212}(t_{ij}) \\ \vdots \\ U_{2L^{U_2}1}(t_{ij}) \\ U_{2L^{U_2}2}(t_{ij}) \end{pmatrix}.$$

The hierarchical and the crossed functional functional random effects differ in the form of $\boldsymbol{z}_i$. Let for simplicity assume that there are two and four levels of the two grouping variables ($L^{U_1} = 2$, $L^{U_2} = 4$), respectively. Then, in total $q = \sum_{g=1}^{G} \rho^{U_g} L^{U_g} = 10$ functional random effects are specified (apart from the smooth error $E_i(t)$).

In the case of hierarchical functional random effects (e.g. four subjects in two groups), the $n \times q$ matrix $\boldsymbol{Z}$ consisting of the $\boldsymbol{z}_i$, $i = 1, \ldots, n$, then has the form

$$\boldsymbol{Z} = \begin{bmatrix} \boldsymbol{z}_1^\top \\ \vdots \\ \boldsymbol{z}_n^\top \end{bmatrix} = \left[\begin{array}{cccc|cccc} 1 & 1 & & & w_1 & & & \\ \vdots & \vdots & & & \vdots & & & \\ 1 & 1 & & & w_{n_1} & & & \\ & & 1 & & & w_{n_1+1} & & \\ & & \vdots & & & \vdots & & \\ & & 1 & & & w_{n_1+n_2} & & \\ 1 & & & 1 & & & w_{n_1+n_2+1} & \\ & 1 & & \vdots & & & \vdots & \\ & \vdots & & 1 & & & w_{n_1+n_2+n_3} & \\ & & 1 & & & & & w_{n_1+n_2+n_3+1} \\ & & \vdots & & & & & \vdots \\ & 1 & & 1 & & & & w_n \end{array}\right],$$

where we implicitly assume for ease of presentation that there are two subjects in each group.

For crossed functional random effects (e.g. two speakers and four words), we assume for simplicity that half of the n_i curves belong to the first and the other to the second level of the first grouping variable. Then, matrix $\boldsymbol{Z}$ is given by

$$\boldsymbol{Z} = \begin{bmatrix} \boldsymbol{z}_1^\top \\ \vdots \\ \boldsymbol{z}_n^\top \end{bmatrix} = \left[\begin{array}{cccccc|cccc} 1 & 1 & & & & & w_1 & & & \\ \vdots & \vdots & & & & & \vdots & & & \\ 1 & 1 & & & & & w_{\frac{n_1}{2}} & & & \\ & & 1 & & & & & w_{\frac{n_1}{2}+1} & & \\ & & \vdots & & & & & \vdots & & \\ & & 1 & & & & & w_{\frac{n_1}{2}+\frac{n_2}{2}} & & \\ & & & 1 & & & & & w_{\frac{n_1}{2}+\frac{n_2}{2}+1} & \\ & & & \vdots & & & & & \vdots & \\ & & & 1 & & & & & w_{\frac{n_1}{2}+\frac{n_2}{2}+\frac{n_3}{2}} & \\ & & & & 1 & & & & & w_{\frac{n_1}{2}+\frac{n_2}{2}+\frac{n_3}{2}+1} \\ & & & & \vdots & & & & & \vdots \\ & & & & 1 & & & & & w_{\frac{n}{2}} \\ & 1 & 1 & & & & w_{\frac{n}{2}+1} & & & \\ & \vdots & \vdots & & & & \vdots & & & \\ & 1 & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}} & & & \\ & & & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}+1} & & \\ & & & \vdots & & & & \vdots & & \\ & & & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}+\frac{n_2}{2}} & & \\ & & & & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}+\frac{n_2}{2}+1} & \\ & & & & \vdots & & & & \vdots & \\ & & & & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}+\frac{n_2}{2}+\frac{n_3}{2}} & \\ & & & & & 1 & & & & w_{\frac{n}{2}+\frac{n_1}{2}+\frac{n_2}{2}+\frac{n_3}{2}+1} \\ & & & & & \vdots & & & & \vdots \\ & & & & & 1 & & & & w_n \end{array}\right].$$

B.3 Form of the constraint matrix

For each $g = 1, \ldots, G$, the constraint matrix $\boldsymbol{W}^{U_g}$ is a block matrix consisting of $\left(\rho^{U_g}\right)^2 \times \frac{\left(\rho^{U_g}\right)^2+1}{2}$ blocks, most of which are zero. The rows and columns of $\boldsymbol{W}^{U_g}$ are sorted as

in matrix $\boldsymbol{M}^{U_g\Delta}$ and in the reduced matrix $\boldsymbol{M}^{U_g\Delta r}$, respectively. The non-zero blocks can be divided into two groups: blocks corresponding to the auto-covariances $K_{ss}^{U_g}(t,t')$, $s=1,\ldots,\rho^{U_g}$, and blocks corresponding to the cross-covariances $K_{ss'}^{U_g}(t,t')$, with $s<s'$.

Let $\left(F_{ss}^{U_g}\right)^2$ denote the number of spline basis functions used for smoothing the auto-covariance $K_{ss}^{U_g}(t,t')$. The blocks for the auto-covariances are of the same form as $\boldsymbol{W}^E$ and given by the $\left(F_{ss}^{U_g}\right)^2 \times F_{ss}^{U_g}(F_{ss}^{U_g}+1)/2$ matrices

$$\begin{bmatrix} \boldsymbol{I}_{\frac{F_{ss}^{U_g}(F_{ss}^{U_g}-1)}{2}} & \boldsymbol{0}_{\frac{F_{ss}^{U_g}(F_{ss}^{U_g}-1)}{2}\times F_{ss}^{U_g}} \\ \boldsymbol{0}_{F_{ss}^{U_g}\times\frac{F_{ss}^{U_g}(F_{ss}^{U_g}-1)}{2}} & \boldsymbol{I}_{F_{ss}^{U_g}} \\ \boldsymbol{I}_{\frac{F_{ss}^{U_g}(F_{ss}^{U_g}-1)}{2}} & \boldsymbol{0}_{\frac{F_{ss}^{U_g}(F_{ss}^{U_g}-1)}{2}\times F_{ss}^{U_g}} \end{bmatrix},$$

where $\boldsymbol{I}_x$ is an identity matrix of dimension x and $\boldsymbol{0}_{x\times y}$ is a null matrix of dimension $x\times y$.

Consider for simplicity the case of bivariate tensor product spline bases, where we can denote $F_{ss',t}^{U_g}$ and $F_{ss',t'}^{U_g}$ the number of marginal spline basis functions for smoothing the cross-covariance $K_{ss'}^{U_g}(t,t')$, $s<s'$, in direction t and t', respectively. Due to the symmetry, we have $F_{ss',t}^{U_g}=F_{s's,t'}^{U_g}$ and $F_{ss',t'}^{U_g}=F_{s's,t}^{U_g}$. Let $F_{ss',b=b'}^{U_g}$ denote the number of coefficients on the diagonal in $\boldsymbol{\Theta}_{ss'}^{U_g}$, which corresponds to the minimum of $F_{ss',t}^{U_g}$ and $F_{ss',t'}^{U_g}$ and denote the number of coefficients below and above the diagonal as $F_{ss',b<b'}^{U_g} \coloneqq \sum_{i=1}^{F_{ss',b=b'}^{U_g}}\left(F_{ss',t'}^{U_g}-i\right)$ and $F_{ss',b>b'}^{U_g} \coloneqq \sum_{i=1}^{F_{ss',b=b'}^{U_g}}\left(F_{ss',t}^{U_g}-i\right)$, respectively. The blocks for the cross-covariances are then $F_{ss',t}^{U_g}F_{ss',t'}^{U_g} \times F_{ss',t}^{U_g}F_{ss',t'}^{U_g}$ diagonal block matrices of the form

$$\begin{bmatrix} \boldsymbol{I}_{F_{ss',b<b'}^{U_g}} & \boldsymbol{0}_{F_{ss',b<b'}^{U_g}\times F_{ss'b=b'}^{U_g}} & \boldsymbol{0}_{F_{ss',b<b'}^{U_g}\times F_{ss',b>b'}^{U_g}} \\ \boldsymbol{0}_{F_{ss',b=b'}^{U_g}\times F_{ss',b<b'}^{U_g}} & \boldsymbol{I}_{F_{ss',b=b'}^{U_g}} & \boldsymbol{0}_{F_{ss',b=b'}^{U_g}\times F_{ss',b>b'}^{U_g}} \\ \boldsymbol{0}_{F_{ss',b>b'}^{U_g}\times F_{ss',b<b'}^{U_g}} & \boldsymbol{0}_{F_{ss',b>b'}^{U_g}\times F_{ss',b=b'}^{U_g}} & \boldsymbol{I}_{F_{ss',b>b'}^{U_g}} \end{bmatrix},$$

when the respective rows correspond to $s<s'$ and $F_{ss',t}^{U_g}F_{ss',t'}^{U_g} \times F_{ss',t}^{U_g}F_{ss',t'}^{U_g}$ anti-diagonal

block matrices of the form

$$\begin{bmatrix} \mathbf{0}_{F^{U_g}_{ss',b<b'} \times F^{U_g}_{ss',b<b'}} & \mathbf{0}_{F^{U_g}_{ss',b<b'} \times F^{U_g}_{ss',b=b'}} & \boldsymbol{I}_{F^{U_g}_{ss',b>b'}} \\ \mathbf{0}_{F^{U_g}_{ss',b=b'} \times F^{U_g}_{ss',b<b'}} & \boldsymbol{I}_{F^{U_g}_{ss',b=b'}} & \mathbf{0}_{F^{U_g}_{ss',b=b'} \times F^{U_g}_{ss',b>b'}} \\ \boldsymbol{I}_{F^{U_g}_{ss',b<b'}} & \mathbf{0}_{F^{U_g}_{ss',b<b'} \times F^{U_g}_{ss',b=b'}} & \mathbf{0}_{F^{U_g}_{ss',b<b'} \times F^{U_g}_{ss',b>b'}} \end{bmatrix},$$

when the respective rows correspond to $s > s'$.

Example with two random effects Consider a grouping variable g with $\rho^{U_g} = 2$ components. Omitting the dimensions of the submatrices for better readability, the constraint matrix $\boldsymbol{W}^{U_g}$ is given by

$$\boldsymbol{W}^{U_g} = \begin{array}{c} \\ (1,1) \\ \\ \\ (1,2) \\ \\ \\ (2,1) \\ \\ \\ (2,2) \\ \\ \end{array} \left[\begin{array}{cc|ccc|cc} \boldsymbol{I} & \mathbf{0} & & & & & \\ \mathbf{0} & \boldsymbol{I} & & & & & \\ \boldsymbol{I} & \mathbf{0} & & & & & \\ \hline & & \boldsymbol{I} & \mathbf{0} & \mathbf{0} & & \\ & & \mathbf{0} & \boldsymbol{I} & \mathbf{0} & & \\ & & \mathbf{0} & \mathbf{0} & \boldsymbol{I} & & \\ \hline & & \mathbf{0} & \mathbf{0} & \boldsymbol{I} & & \\ & & \mathbf{0} & \boldsymbol{I} & \mathbf{0} & & \\ & & \boldsymbol{I} & \mathbf{0} & \mathbf{0} & & \\ \hline & & & & & \boldsymbol{I} & \mathbf{0} \\ & & & & & \mathbf{0} & \boldsymbol{I} \\ & & & & & \boldsymbol{I} & \mathbf{0} \end{array}\right],$$

yielding the reduced design matrix $\boldsymbol{M}^{U_g \Delta r}$

$$\left[\underbrace{\boldsymbol{M}^{U_g\Delta}_{11,b<b'} + \boldsymbol{M}^{U_g\Delta}_{11,b>b'} | \boldsymbol{M}^{U_g\Delta}_{11,b=b'}}_{s=s'=1} \mid \underbrace{\boldsymbol{M}^{U_g\Delta}_{12,b<b'} + \boldsymbol{M}^{U_g\Delta}_{21,b>b'} | \boldsymbol{M}^{U_g\Delta}_{12,b=b'} + \boldsymbol{M}^{U_g\Delta}_{21,b=b'} | \boldsymbol{M}^{U_g\Delta}_{12,b>b'} + \boldsymbol{M}^{U_g\Delta}_{21,b<b'}}_{s<s'(s=1,s'=2)} \mid \underbrace{\boldsymbol{M}^{U_g\Delta}_{22,b<b'} + \boldsymbol{M}^{U_g\Delta}_{22,b>b'} | \boldsymbol{M}^{U_g\Delta}_{22,b=b'}}_{s=s'=2}\right].$$

Example with three random effects Analogously, for a grouping variable g with $\rho^{U_g} = 3$ components the constraint matrix $\boldsymbol{W}^{U_g}$ is given by

$$\boldsymbol{W}^{U_g} = \begin{array}{c} \\ (1,1) \\ (1,2) \\ (1,3) \\ (2,1) \\ (2,2) \\ (2,3) \\ (3,1) \\ (3,2) \\ (3,3) \end{array} \left[\begin{array}{cc|ccc|ccc|cc|ccc|cc} \boldsymbol{I} & \boldsymbol{0} & & & & & & & & & & & & & \\ \boldsymbol{0} & \boldsymbol{I} & & & & & & & & & & & & & \\ \boldsymbol{I} & \boldsymbol{0} & & & & & & & & & & & & & \\ \hline & & \boldsymbol{I} & \boldsymbol{0} & \boldsymbol{0} & & & & & & & & & & \\ & & \boldsymbol{0} & \boldsymbol{I} & \boldsymbol{0} & & & & & & & & & & \\ & & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{I} & & & & & & & & & & \\ \hline & & & & & \boldsymbol{I} & \boldsymbol{0} & \boldsymbol{0} & & & & & & & \\ & & & & & \boldsymbol{0} & \boldsymbol{I} & \boldsymbol{0} & & & & & & & \\ & & & & & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{I} & & & & & & & \\ \hline & & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{I} & & & & & & & & & & \\ & & \boldsymbol{0} & \boldsymbol{I} & \boldsymbol{0} & & & & & & & & & & \\ & & \boldsymbol{I} & \boldsymbol{0} & \boldsymbol{0} & & & & & & & & & & \\ \hline & & & & & \boldsymbol{I} & \boldsymbol{0} & & & & & & & & \\ & & & & & \boldsymbol{0} & \boldsymbol{I} & & & & & & & & \\ & & & & & \boldsymbol{I} & \boldsymbol{0} & & & & & & & & \\ \hline & & & & & & & & \boldsymbol{I} & \boldsymbol{0} & \boldsymbol{0} & & & & \\ & & & & & & & & \boldsymbol{0} & \boldsymbol{I} & \boldsymbol{0} & & & & \\ & & & & & & & & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{I} & & & & \\ \hline & & & & & & & & & & & \boldsymbol{I} & \boldsymbol{0} & & \\ & & & & & & & & & & & \boldsymbol{0} & \boldsymbol{I} & & \\ & & & & & & & & & & & \boldsymbol{I} & \boldsymbol{0} & & \\ \hline & & & & & & & & & & & & & \boldsymbol{I} & \boldsymbol{0} \\ & & & & & & & & & & & & & \boldsymbol{0} & \boldsymbol{I} \\ & & & & & & & & & & & & & \boldsymbol{I} & \boldsymbol{0} \end{array}\right].$$

C Supplementary application details and results

C.1 CD4 cell count data

In AIDS research, the CD4 cell counts as a function of time since seroconversion (SC) – the time at which HIV becomes detectable – often serve as a longitudinally measured biomarker which provides insight into the progression of the disease. As the virus destroys the CD4 cells, a decreasing number of CD4 cells indicates a progress of the disease. The considered data set is part of the Multicenter AIDS Cohort Study (MACS; Kaslow et al., 1987). It contains the CD4 cell count trajectories of 366 HIV infected subjects collected from month -18 to month 42 since SC. Measurements were taken at roughly semi-annual visits yielding a total of 1888 CD4 cell counts per milliliter of blood, with between 1 to 11 counts per subject and a median of 5. To reduce skewness, we base our analysis on the square root of the CD4 cell counts, which are depicted in Figure 6, with some trajectories highlighted for better display and an estimated overall mean function. We can see that on average, the CD4 cell counts are decreasing over time. The data are available in R-package refund

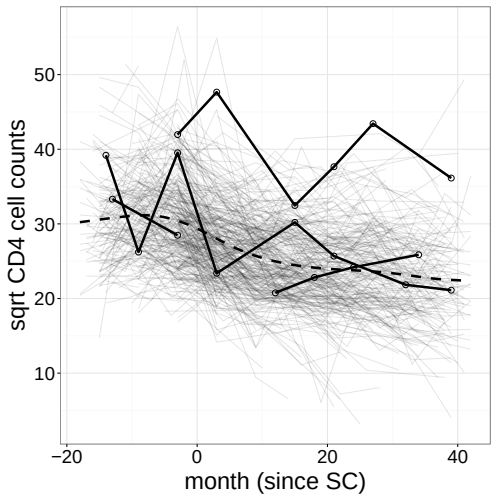

Figure 6: Square root of observed CD4 cell count trajectories plotted against the months since SC. Shown are the trajectories of 366 HIV infected subjects. Some trajectories are highlighted for better display and an estimated smooth mean function (dashed) is shown.

(Huang et al., 2016) and are further described in Goldsmith et al. (2013). Similar data from this study were previously analyzed in e.g. Diggle et al. (2002), Yao et al. (2005), and Peng and Paul (2009).

We fit Model (2.1) with only one fRI for each curve and an overall mean $\mu(t)$. In order to predict the continuous subject-specific trajectories with only few observations per subject available, we perform an FPCA based on our fast covariance smoothing approach (TRI-CONSTR and TRI-CONSTR-W with weights on the diagonal cross products). We demonstrate the similarity to the computationally less efficient approach proposed in Cederbaum et al. (2016) (WHOLE), in which all cross products enter the estimation. Moreover, we show that boundary effects occur on the diagonal when only the triangular surface is estimated without a symmetry constraint (TRI). We compare our results to those obtained from applying the covariance smoothing approach proposed by Xiao et al. (2016b) (FACE) using R-function face.sparse in package face (Xiao et al., 2016a). Moreover, we compare with FACE-STEP-1, a modification of FACE, in which the covariance of the cross products is not accounted for and thus only the first step of the three-step procedure is performed.

We use 13 cubic B-spline basis functions for the estimation of the mean function and as marginal bases for the estimation of the auto-covariance surface using tensor products. To avoid over-fitting, we add a second order difference penalty. For our approach, we use the Kronecker sum penalty (cf. Section 5). We use equidistant knots in function face.sparse instead of the default (quantile based knots) which would require an adapted penalty that is not implemented. The equidistant grid, on which the mean and the auto-covariance are evaluated, is of length $D = 100$, with values between -18 and 42. Note that so far function face.sparse assumes that the function argument takes values in the unit interval. We thus transformed the function argument and re-transformed the results to the original interval $[-18, 42]$ after the estimation. Note that in order to ensure orthonormality with respect to the L^2-scalar product, we rescale the eigenfunctions and accordingly the eigenvalues after re-transforming the function argument. We truncate the number of eigenfunctions using a pre-specified proportion of explained variance of $L = 0.99$. Note that we use the proportion of explained variance in the observed trajectories, whereas Xiao et al. (2016a) use that in $E_i(t)$.

Figure 7 and Figure 8 show the estimated covariance surfaces, reconstructed after trun-

cation from the estimated eigenvalues and eigenfunctions, for our approach, WHOLE, and TRI and for FACE and FACE-STEP-1, respectively. In the bottom of the two figures, we also depict the truncated estimated eigenfunctions. Table 1 additionally gives the truncated estimated eigenvalues and the estimated error variance. As in the application to the phonetics data in Section 6.2, we obtain the same number of eigenfunctions (two) for our approach and WHOLE and a higher number (eleven) of wigglier eigenfunctions for TRI. FACE-STEP-1 also yields two eigenfunctions and FACE yields four eigenfunctions. We can see small differences in the resulting covariance surfaces of our approaches TRI-CONSTR and TRI-CONSTR-W. As expected, the latter is slightly more similar to that of WHOLE. The estimated surface of FACE-STEP-1 is also similar but a little smoother, whereas the the estimated surface of FACE is less smooth. As expected, the estimated surface of TRI shows a clear difference to the others on the diagonal, where it is much wigglier. The interpretation of the first and second eigenfunction is similar for all compared methods. The first eigenfunction is almost a vertical shift and thus describes the level of the (square root) CD4 cell counts. HIV infected individuals with negative basis weights for the first component tend to have a higher number of CD4 cells during the whole time interval [-18,42] than individuals with positive basis weights. The second eigenfunction gives insight in how fast the disease progresses. Individuals with negative basis weights for the second eigenfunction tend to have a faster decrease in CD4 cells than individuals with positive basis weights.

Table 1: Truncated estimated eigenvalues and estimated error variance for all compared methods.

	$\hat{\nu}_1$	$\hat{\nu}_2$	$\hat{\nu}_3$	$\hat{\nu}_4$	$\hat{\nu}_5$	$\hat{\nu}_6$	$\hat{\nu}_7$	$\hat{\nu}_8$	$\hat{\nu}_9$	$\hat{\nu}_{10}$	$\hat{\nu}_{11}$	$\hat{\sigma}^2$
TRI-CONSTR	1170.37	184.73										15.54
TRI-CONSTR-W	1173.96	178.71										15.63
WHOLE	1174.96	178.03										15.57
TRI	1191.80	205.31	61.88	26.95	17.21	9.47	6.57	4.45	3.55	2.75	2.30	12.13
FACE	1162.51	280.53	22.28	12.87								13.70
FACE-STEP-1	1161.84	191.53										15.45

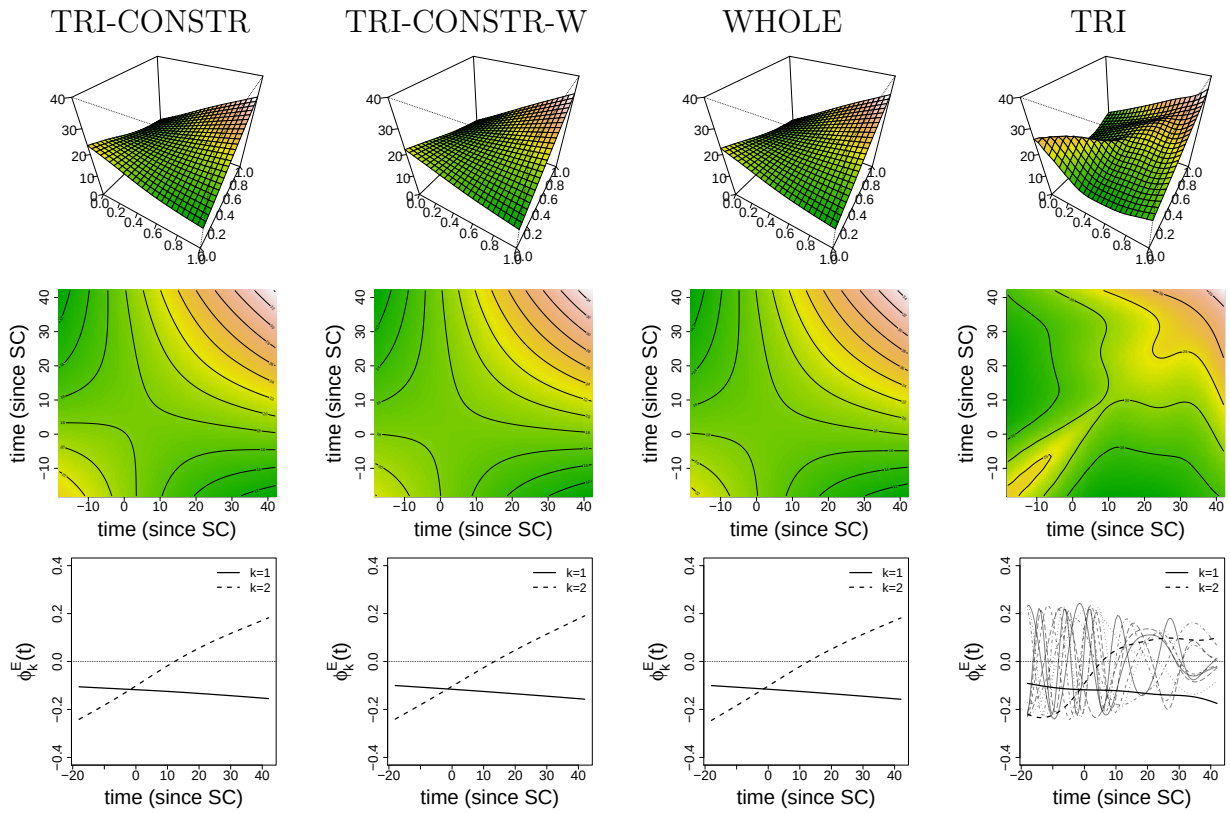

Figure 7: Results for the curve-specific fRI for TRI-CONSTR, TRI-CONSTR-W, WHOLE, and TRI. Top row: estimated covariance surfaces. Middle row: contours of the estimated covariance surfaces. Bottom row: estimated corresponding eigenfunctions.

C.2 Phonetics data

In the following, we show additional results for our application to the phonetics data (cf. Section 6.2), including the estimated auto-covariances for the smooth error $E_i(t)$, the estimated eigenvalues for both random processes, and the estimated error variance.

Figure 9 depicts the estimated surfaces and contours of the auto-covariance of the smooth error $E_i(t)$, reconstructed after truncation from the estimated eigenvalues and eigenfunctions which are shown in the bottom of the figure. As for the auto-covariance of the fRI for speakers, we can see from Figure 9 that the two estimates based on our symmetric smoother

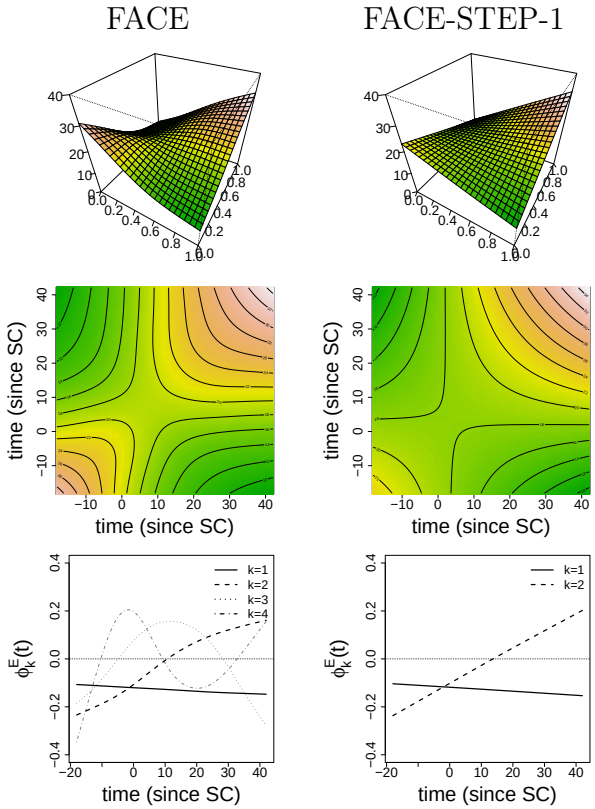

Figure 8: Results for the curve-specific fRI for FACE and FACE-STEP-1. Top row: estimated covariance surfaces. Middle row: contours of the estimated covariance surfaces. Bottom row: estimated corresponding eigenfunctions.

(TRI-CONSTR, TRI-CONSTR-W) are very similar to each other and to the one obtained by using all cross products (WHOLE). Again, we obtain wigglier estimates for TRI – especially on the diagonal of the estimated surface. This corresponds to the fact that for TRI, the error variance is estimated to be zero. The first three eigenfunctions are very similar for all four compared methods. For TRI, however, nine more (high-frequency) eigenfunctions are chosen, yielding a wigglier surface estimate.

Table 2 gives the complete variance decomposition for our model. The upper table shows the truncated estimated eigenvalues of $\hat{K}^B$ for the four compared methods. The lower table shows the truncated estimated eigenvalues of the smooth error as well as the estimated

error variance. For better display, all values are multiplied with 10^3. It shows that the first two [three] estimated eigenvalues for $K^B(t,t')$ $[K^E(t,t')]$ are very similar for all smoothing methods and that TRI-CONSTR-W and WHOLE are most similar. The estimated error variance is slightly higher for TRI-CONSTR than for the other approaches.

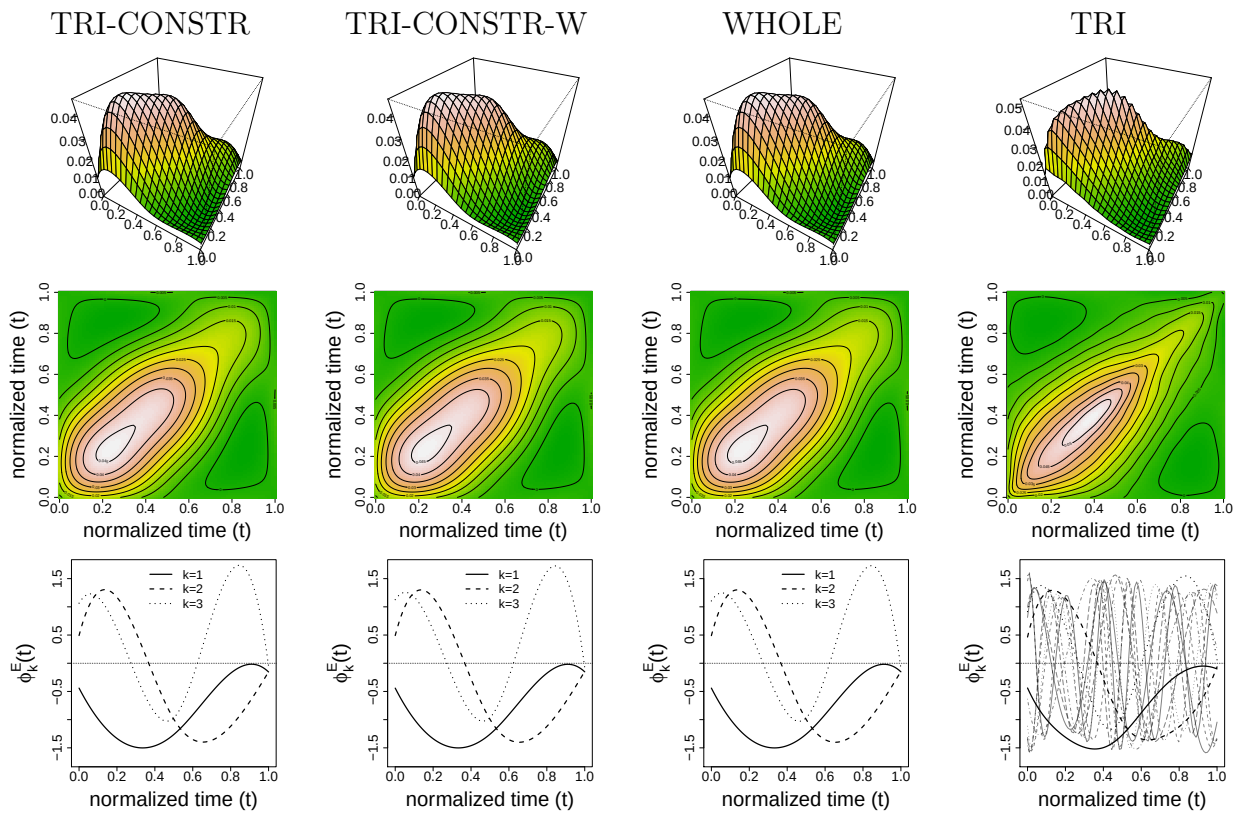

Figure 9: Results for the smooth error curve $E_i(t)$ using the four smoothing methods. Top row: estimated covariance surfaces. Middle row: contours of the estimated covariance surfaces. Bottom row: estimated corresponding eigenfunctions $\phi_k^E(t)$.

Table 2: Truncated estimated eigenvalues of $K^B(t,t')$, $\hat{\nu}_k^B \cdot 10^3$, and of $K^E(t,t')$, $\hat{\nu}_k^E \cdot 10^3$, and estimated error variance $\hat{\sigma}^2 \cdot 10^3$ for all compared methods.

	$\hat{\nu}_1^B$	$\hat{\nu}_2^B$	$\hat{\nu}_3^B$	$\hat{\nu}_4^B$
TRI-CONSTR	5.84	3.24		
TRI-CONSTR-W	5.84	3.23		
WHOLE	5.84	3.23		
TRI	5.85	3.27	0.42	0.24

	$\hat{\nu}_1^E$	$\hat{\nu}_2^E$	$\hat{\nu}_3^E$	$\hat{\nu}_4^E$	$\hat{\nu}_5^E$	$\hat{\nu}_6^E$	$\hat{\nu}_7^E$	$\hat{\nu}_8^E$	$\hat{\nu}_9^E$	$\hat{\nu}_{10}^E$	$\hat{\nu}_{11}^E$	$\hat{\nu}_{12}^E$	$\hat{\sigma}^2$
TRI-CONSTR	19.52	7.56	2.74										4.15
TRI-CONSTR-W	19.53	7.59	2.73										3.97
WHOLE	19.53	7.59	2.73										3.94
TRI	19.67	7.77	2.95	1.37	0.93	0.63	0.50	0.36	0.30	0.23	0.20	0.16	0.00

D Supplementary simulation details and results

D.1 Generation Details

For the scenario with crossed fRIs, we use two, one, and three eigenfunctions, estimated from the phonetics data, for the generation of the auto-covariances of processes $B_i(t)$, $C_i(t)$, and $E_i(t)$, respectively. The resulting auto-covariance surfaces are shown in Figure 10. The eigenfunctions used for data generation are shown in the bottom of the Figure. The corresponding eigenvalues used for data generation are shown in Table 3, where also the error variance is given. The values are multiplied with 10^3 for better display. The underlying mean function is depicted in Figure 11.

Table 3: Eigenvalues $\nu_k^X \cdot 10^3$, $X \in \{B, C, E\}$, and error variance $\sigma^2 \cdot 10^3$ used for data generation for the scenario with crossed fRIs.

ν_1^B	ν_2^B	ν_1^C	ν_1^E	ν_2^E	ν_3^E	σ^2
5.86	2.71	8.89	19.05	7.53	2.66	5.62

D.2 Measures of goodness of fit

We use root relative mean squared errors (rrMSEs) as measures of goodness of fit for all model components (cp. Cederbaum et al., 2016).

For vector-valued estimates $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_L)^\top$, we define the rrMSE as

$$\text{rrMSE}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}) = \sqrt{\frac{\frac{1}{L}\sum_{l=1}^{L}\left(\theta_l - \hat{\theta}_l\right)^2}{\frac{1}{L}\sum_{l=1}^{L}\theta_l{}^2}}. \tag{D.1}$$

We use that for the random basis weights ξ_{lk}^X, (D.1) is approximately $\sqrt{{}^1/_{L^X}\sum_{l=1}^{L^X}\left(\xi_{lk}^X - \hat{\xi}_{lk}^X\right)^2 / \nu_k^X}$, $X \in \{B, C, E\}$.

The form of the rrMSE for scalar estimates results as special case of D.1 with $L = 1$.

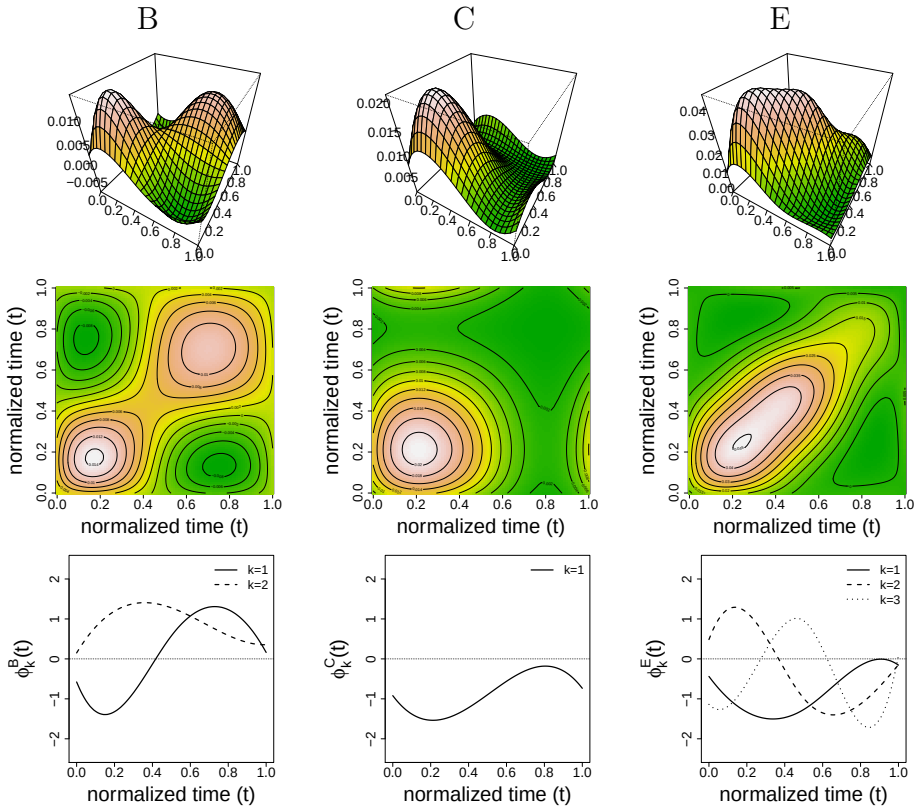

Figure 10: Auto-covariances $K^B(t,t')$, $K^C(t,t')$, and $K^E(t,t')$ and their eigenfunctions used for the data generation for the scenario with crossed fRIs. Top row: covariance surfaces. Middle row: contours of the covariance surfaces. Bottom row: corresponding eigenfunctions $\phi_k^B(t)$, $\phi_k^C(t)$, and $\phi_k^E(t)$.

For all functions $\theta(t)$, we approximate the integrals by sums and obtain

$$\text{rrMSE}\left[\theta(\cdot), \hat{\theta}(\cdot)\right] = \sqrt{\frac{\frac{1}{D}\sum_{d=1}^{D}\left[\theta(t_d) - \hat{\theta}(t_d)\right]^2}{\frac{1}{D}\sum_{d=1}^{D}\theta(t_d)^2}}. \tag{D.2}$$

As the eigenfunctions are only unique up to sign, we also compute the rrMSEs of the estimated eigenfunctions mirrored around the x-axis and choose the smaller rrMSE. For the

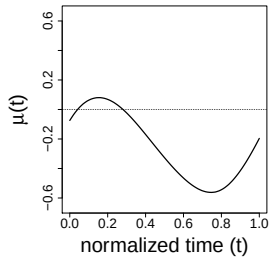

Figure 11: Mean function used for the data generation for the scenario with crossed fRIs.

random processes, we additionally average over the respective levels. For centered processes, we use that the denominator simplifies to the average variance.

For bivariate functions, such as the auto-covariances, we define

$$\text{rrMSE}\left[\theta(\cdot,\cdot),\hat{\theta}(\cdot,\cdot)\right] = \sqrt{\frac{\frac{1}{D^2}\sum_{t_d,t_{d'}=1}^{D}\left(\theta(t_d,t_{d'})-\hat{\theta}(t_d,t_{d'})\right)^2}{\frac{1}{D^2}\sum_{t_d,t_{d'}=1}^{D}\theta(t_d,t_{d'})^2}}. \tag{D.3}$$

D.3 Results for the scenario with independent curves

In the following, we show the complete results for the remaining ten settings (Setting 2 – Setting 11) for the scenario with independent curves. Table 4 lists the different settings we consider for this scenario. In Figure 12 to Figure 21, we depict boxplots of the rrMSEs based on the 200 simulation runs for all model components.

Table 4: Specification of the eleven considered settings for the scenario with independent curves. The simple eigenfunctions are given as $\{\phi_1(t) = 1, \phi_2(t) = \sqrt{3}\,(2t-1)\}$, the complex eigenfunctions are given as $\{\phi_1(t) = \sin(2\pi t), \phi_2(t) = \cos(2\pi t)\}$. The results for Setting 1 are shown in Section 7.2.

Setting	grid	eigenfunctions	eigenvalues	error variance
Setting 1	dense	complex	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.05$
Setting 2	dense	complex	$\nu_1 = 2$, $\nu_2 = 1$,	$\sigma^2 = 0.5$
Setting 3	dense	simple	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.05$
Setting 4	dense	simple	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.5$
Setting 5	sparse	simple	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.05$
Setting 6	sparse	simple	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.5$
Setting 7	dense	complex	$\nu_1 = 0.15$, $\nu_2 = 0.075$	$\sigma^2 = 0.05$
Setting 8	dense	complex	$\nu_1 = 0.15$, $\nu_2 = 0.075$	$\sigma^2 = 0.5$
Setting 9	dense	simple	$\nu_1 = 0.15$, $\nu_2 = 0.075$	$\sigma^2 = 0.05$
Setting 10	dense	simple	$\nu_1 = 0.15$, $\nu_2 = 0.075$	$\sigma^2 = 0.5$
Setting 11	dense	complex	$\nu_1 = 2$, $\nu_2 = 1$	$\sigma^2 = 0.01$

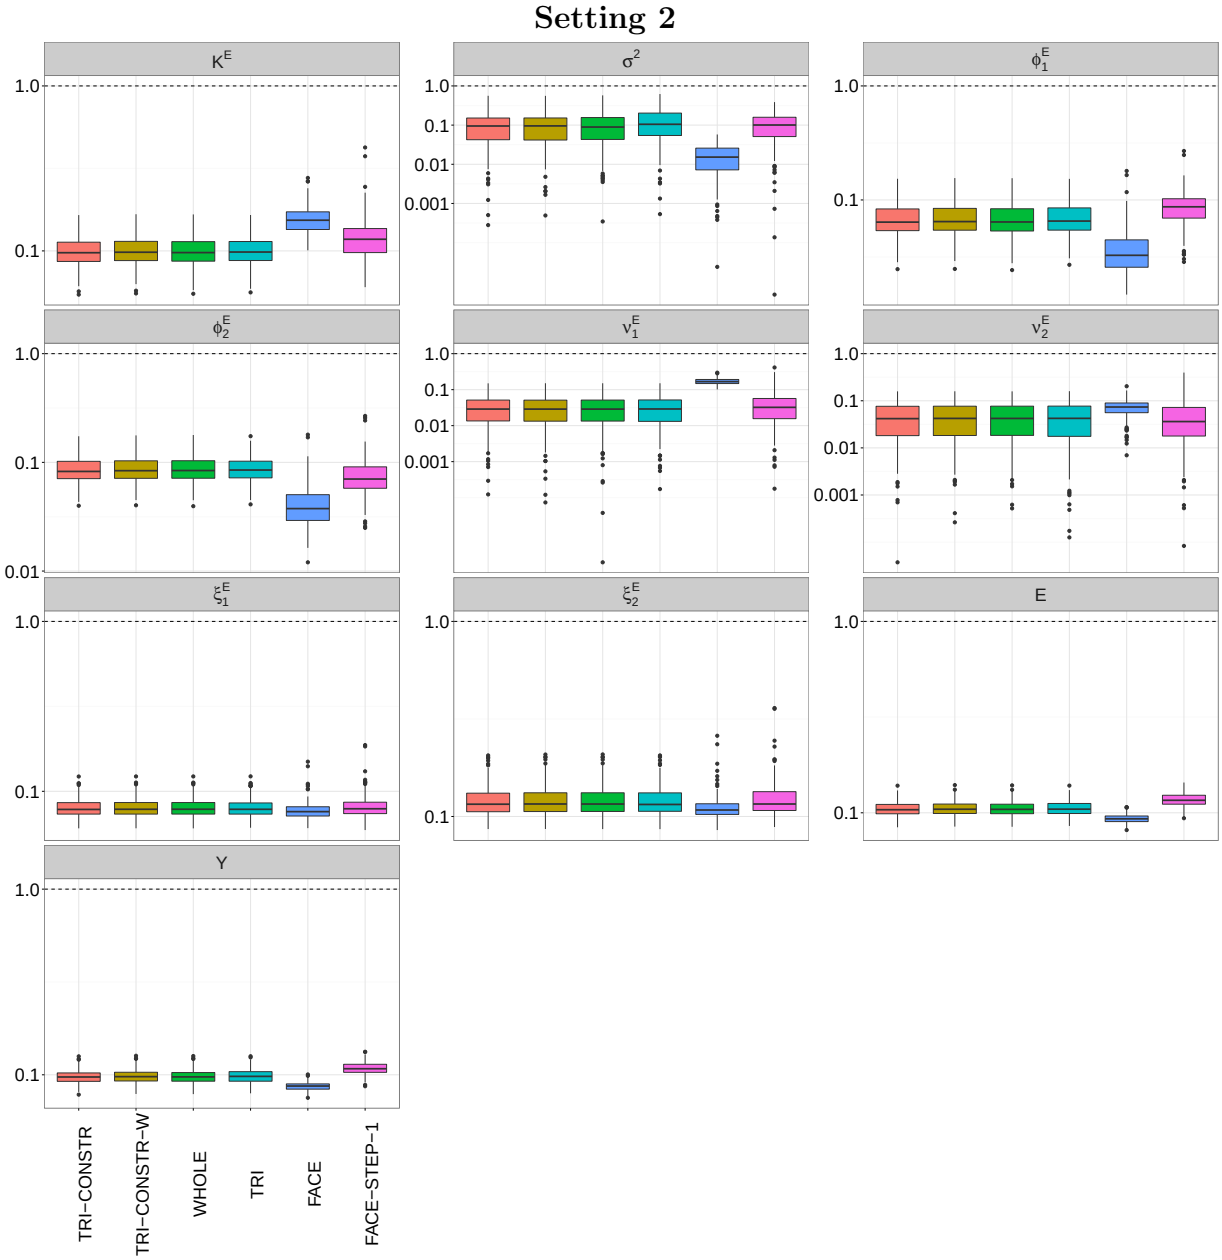

Figure 12: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 2 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

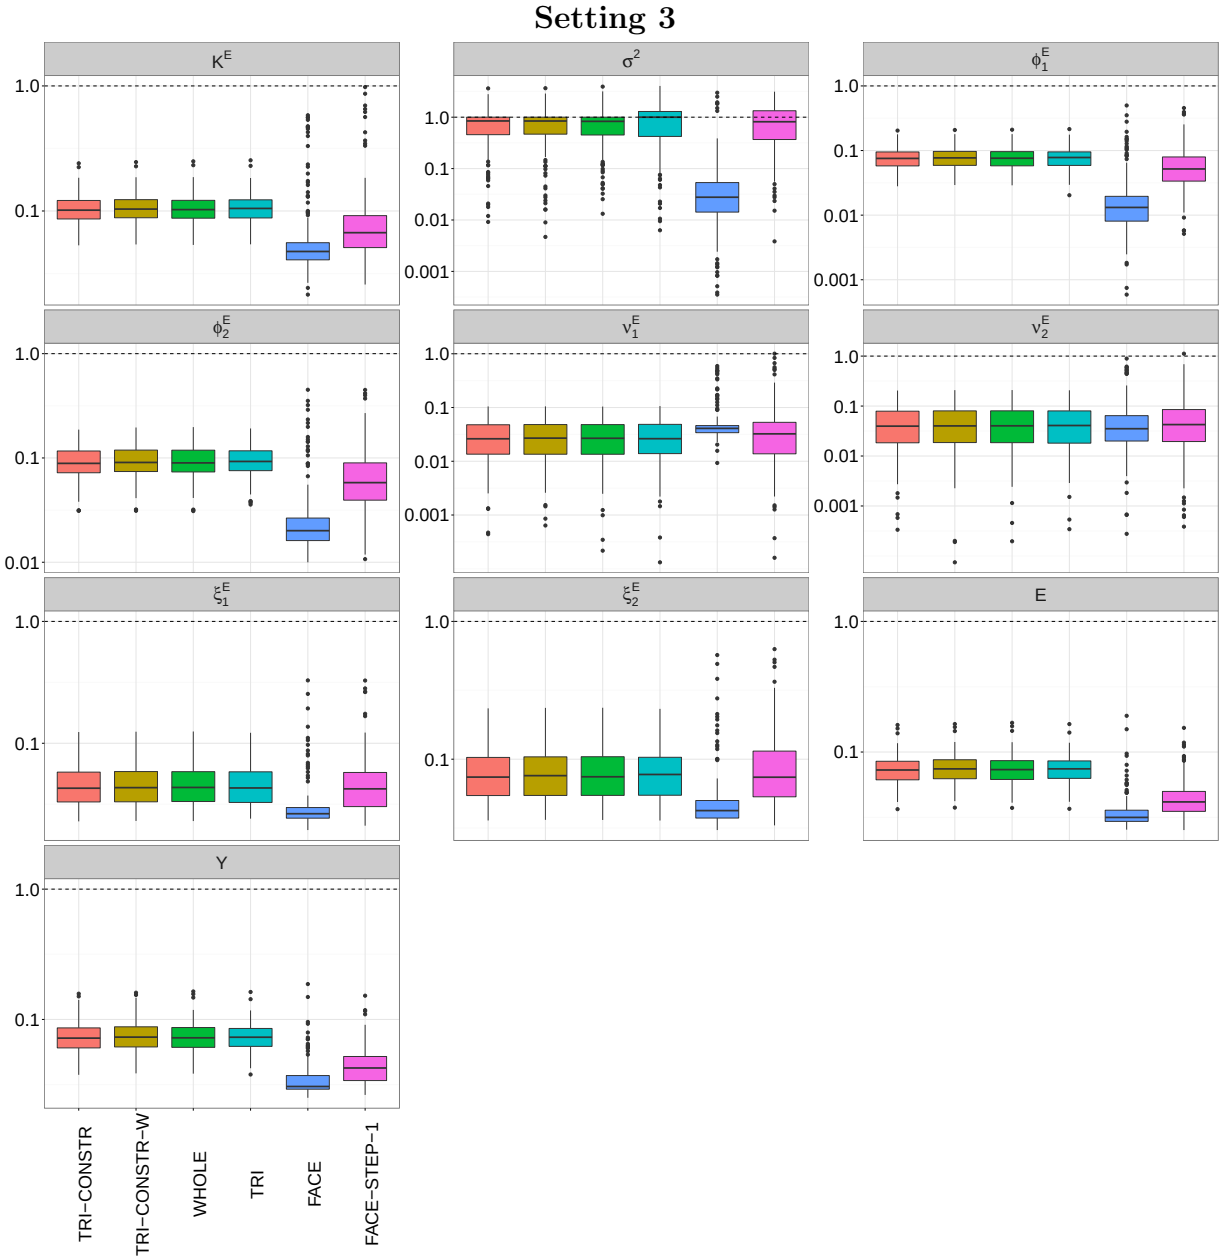

Figure 13: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 3 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

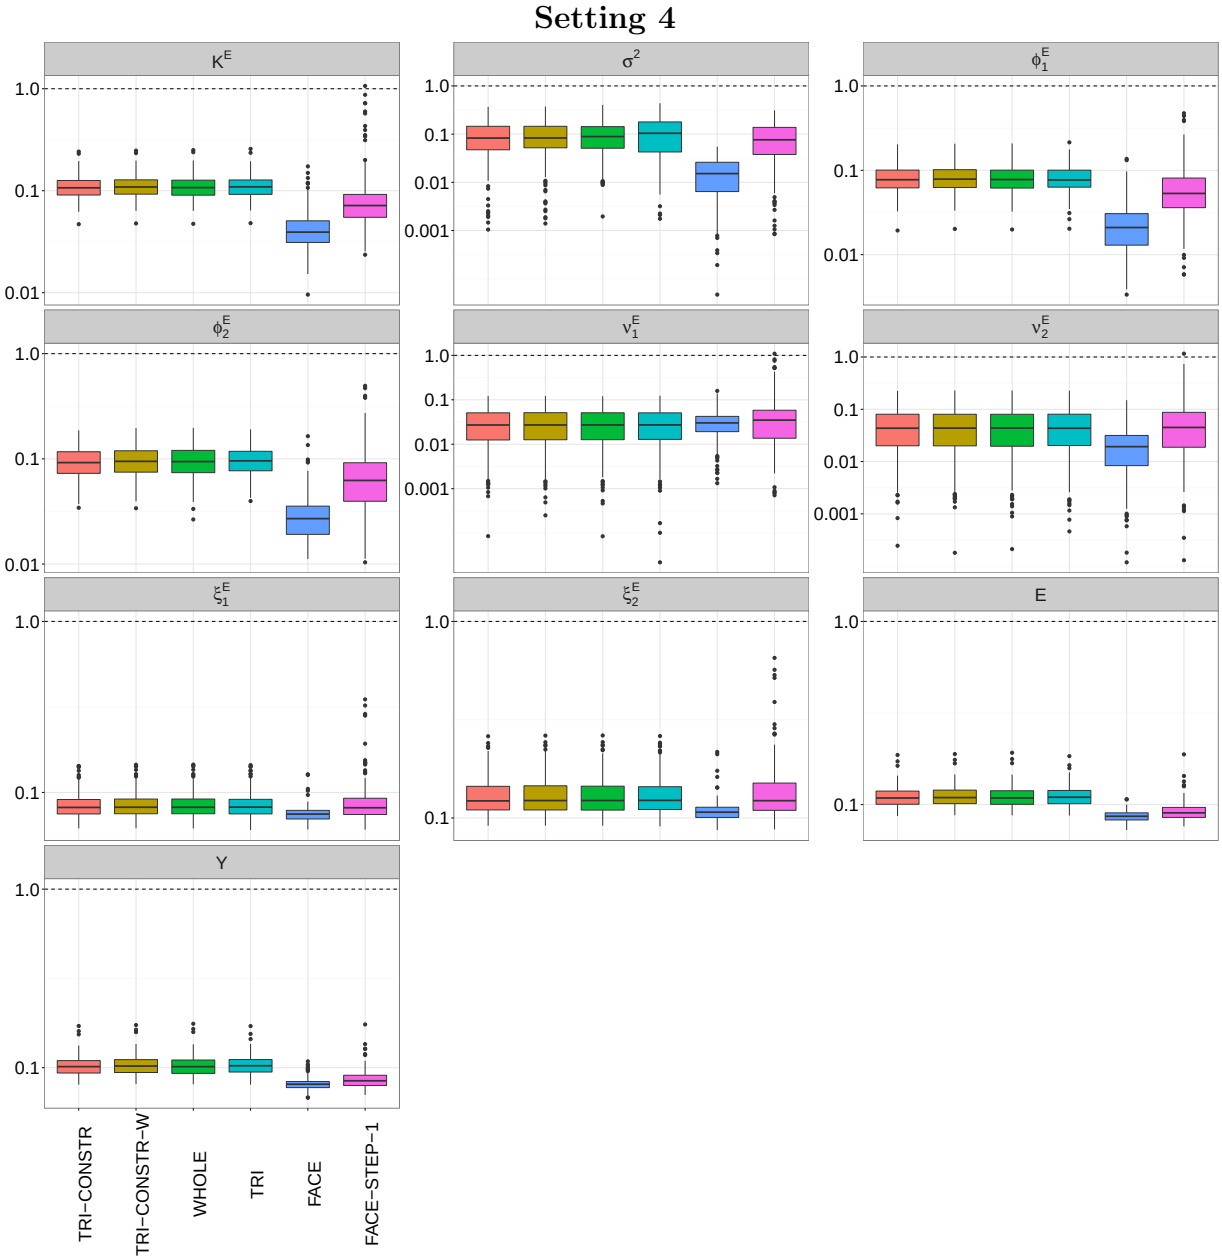

Figure 14: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 4 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

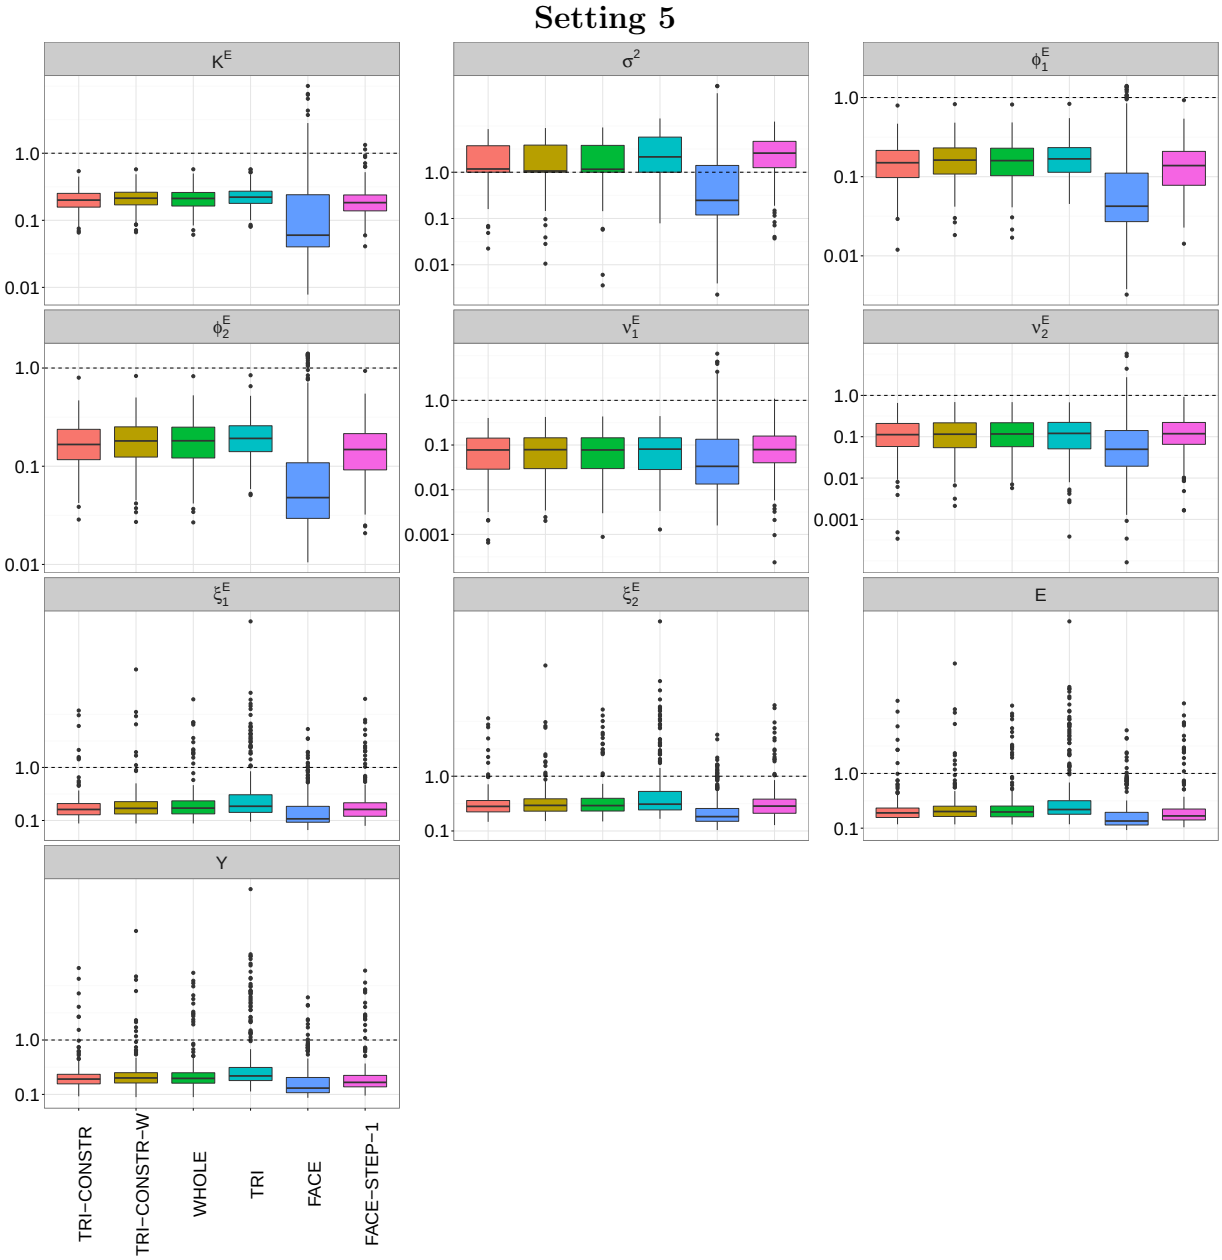

Figure 15: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 5 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

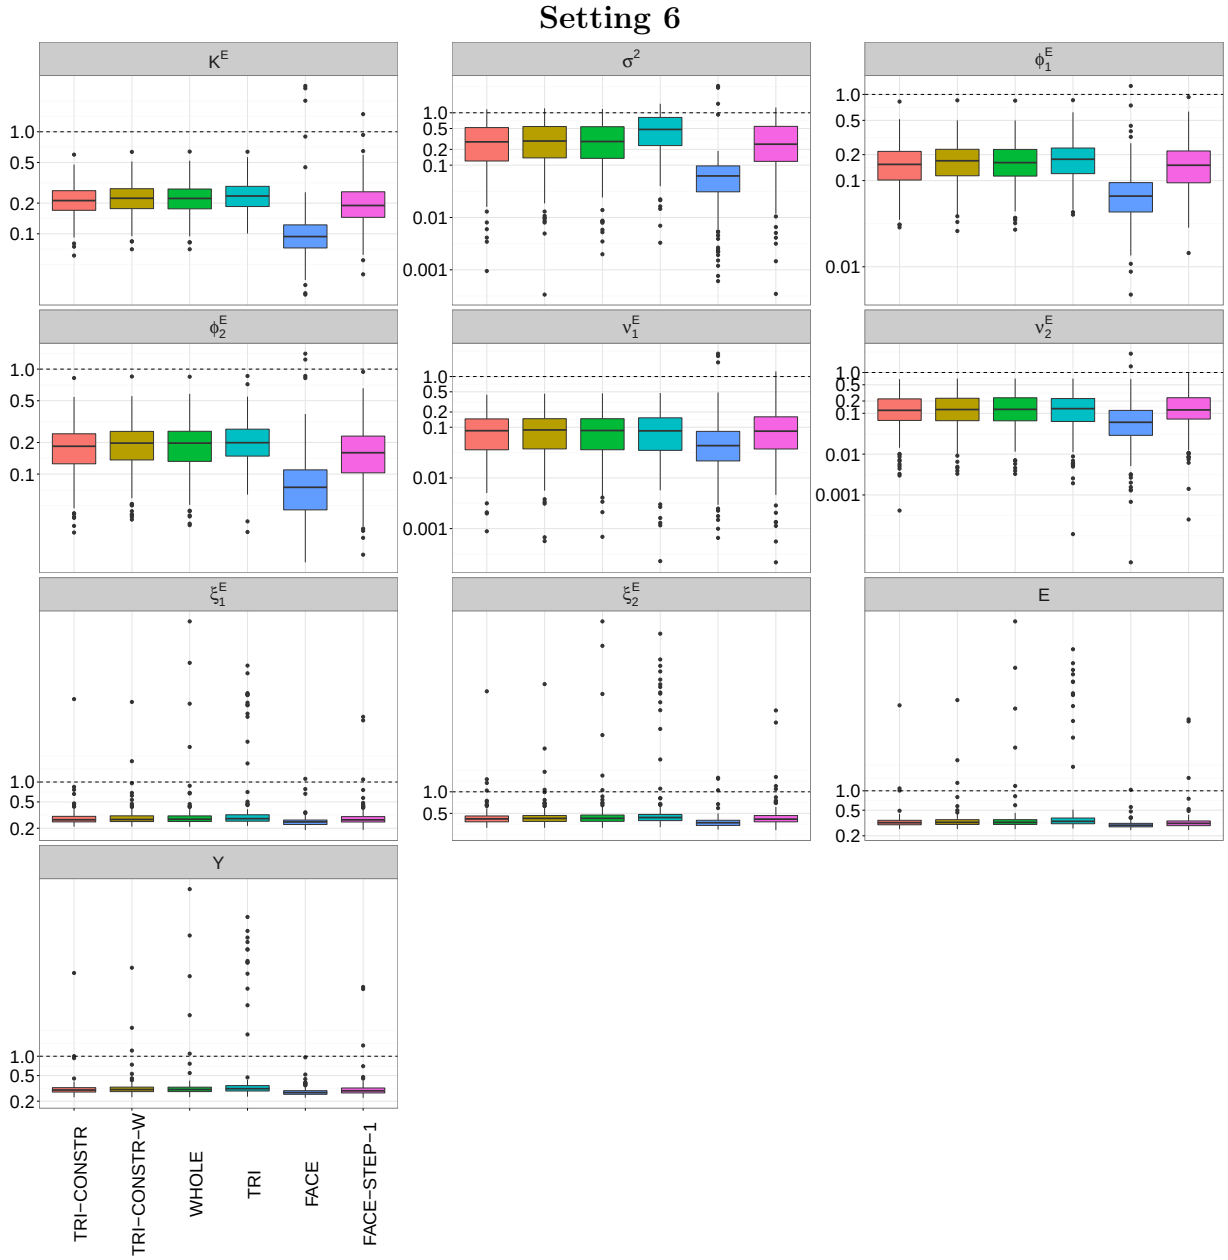

Figure 16: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 6 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

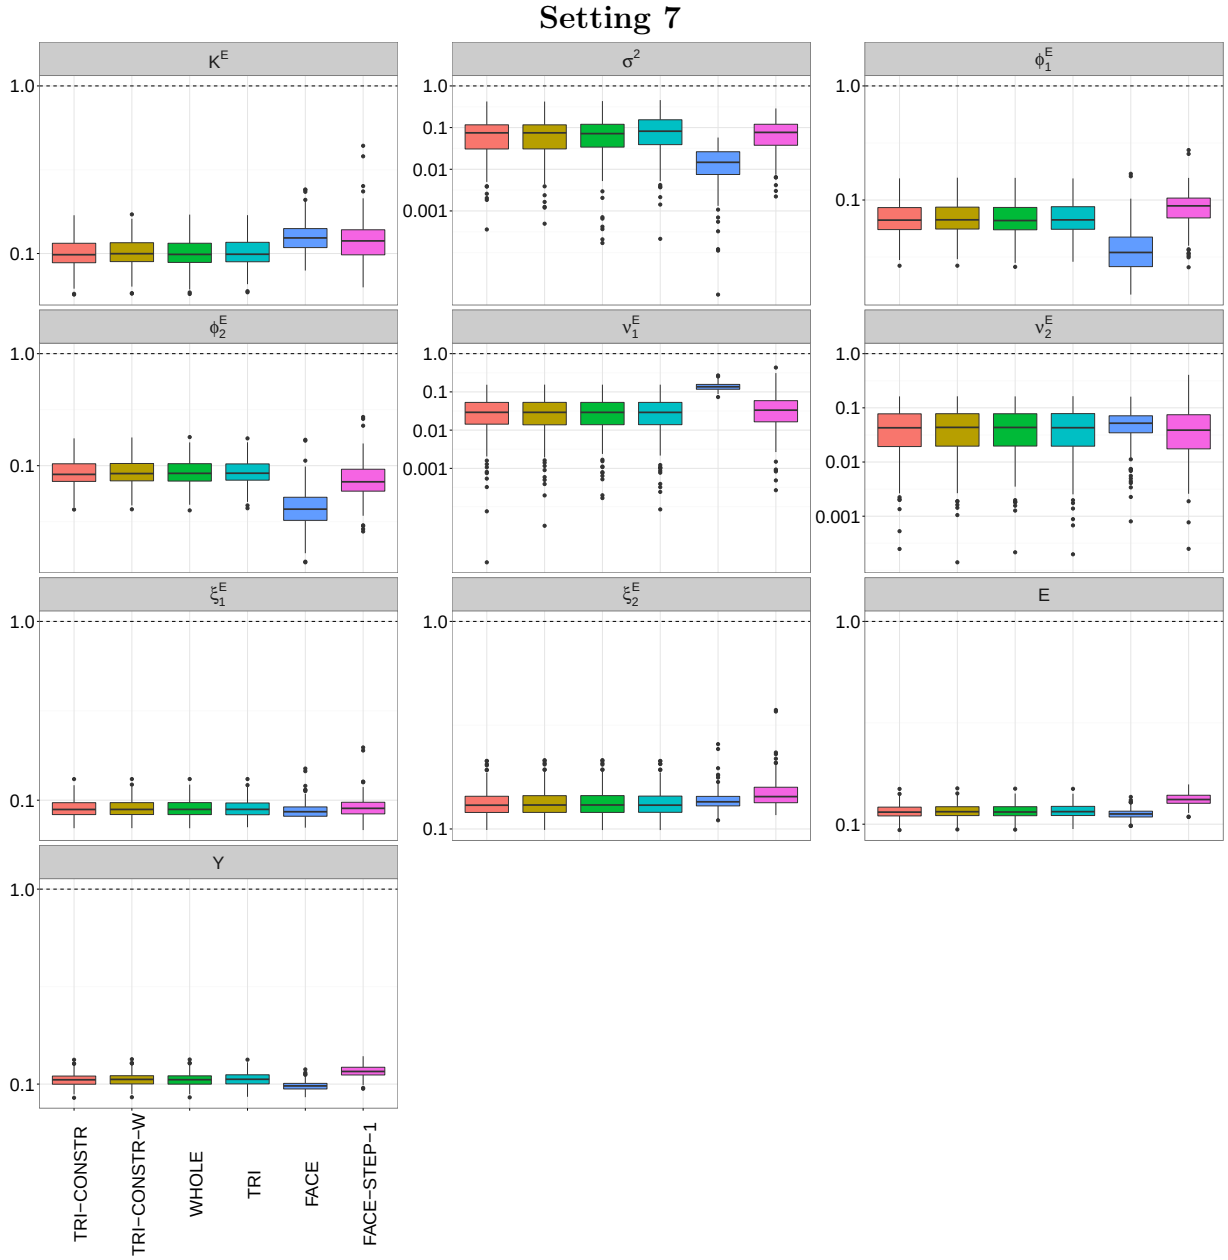

Figure 17: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 7 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

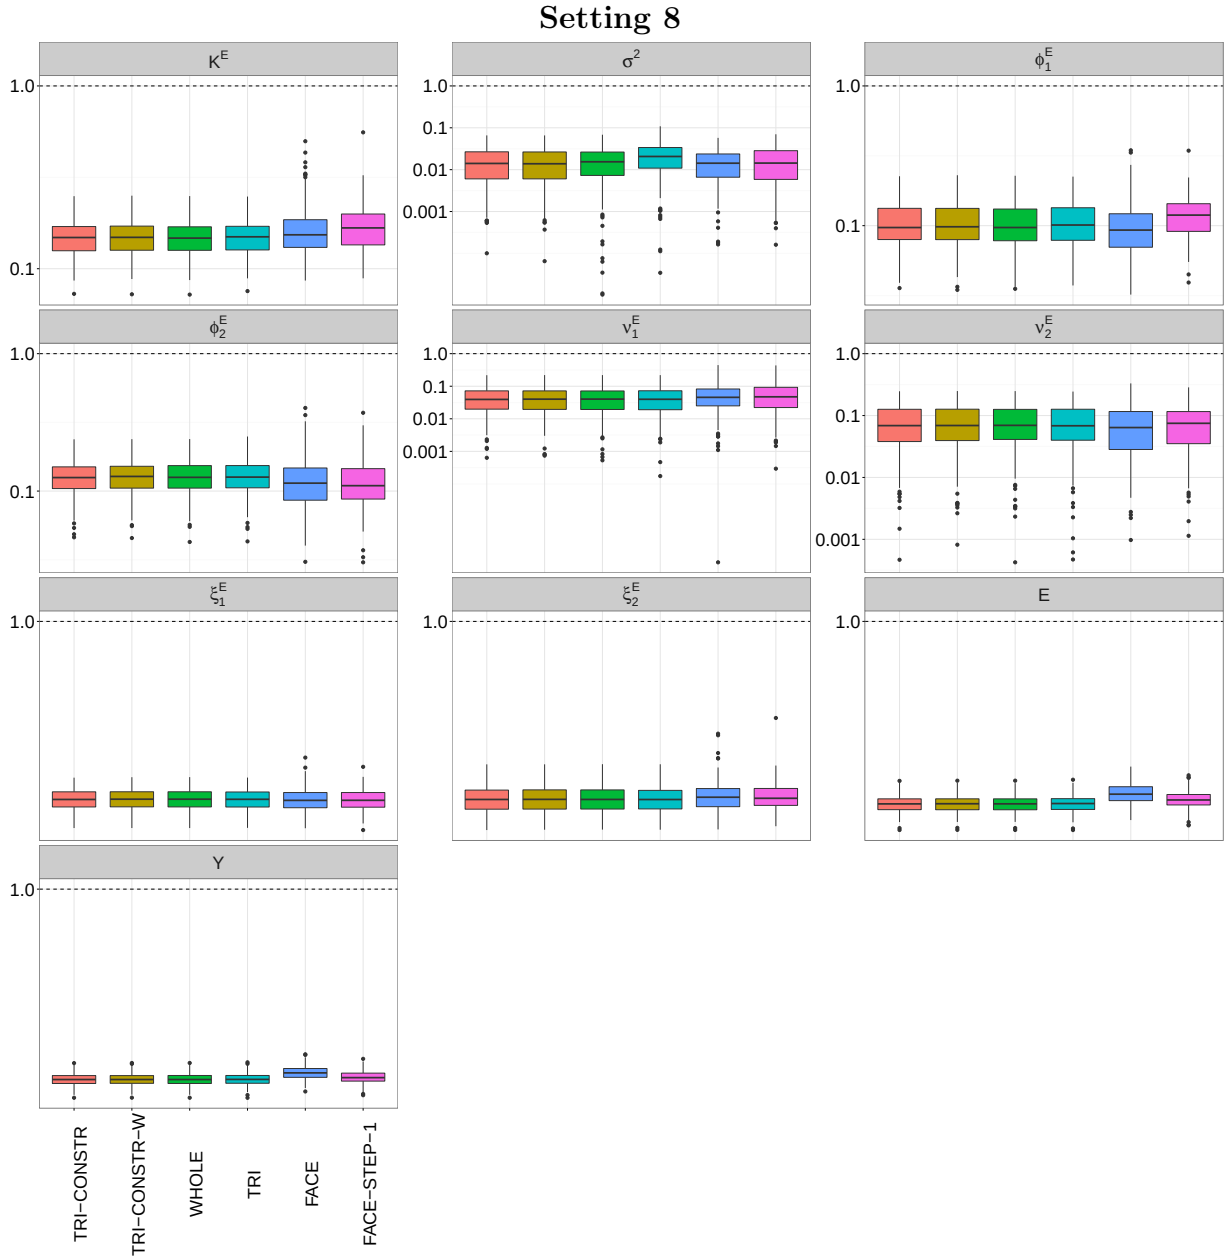

Figure 18: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 8 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

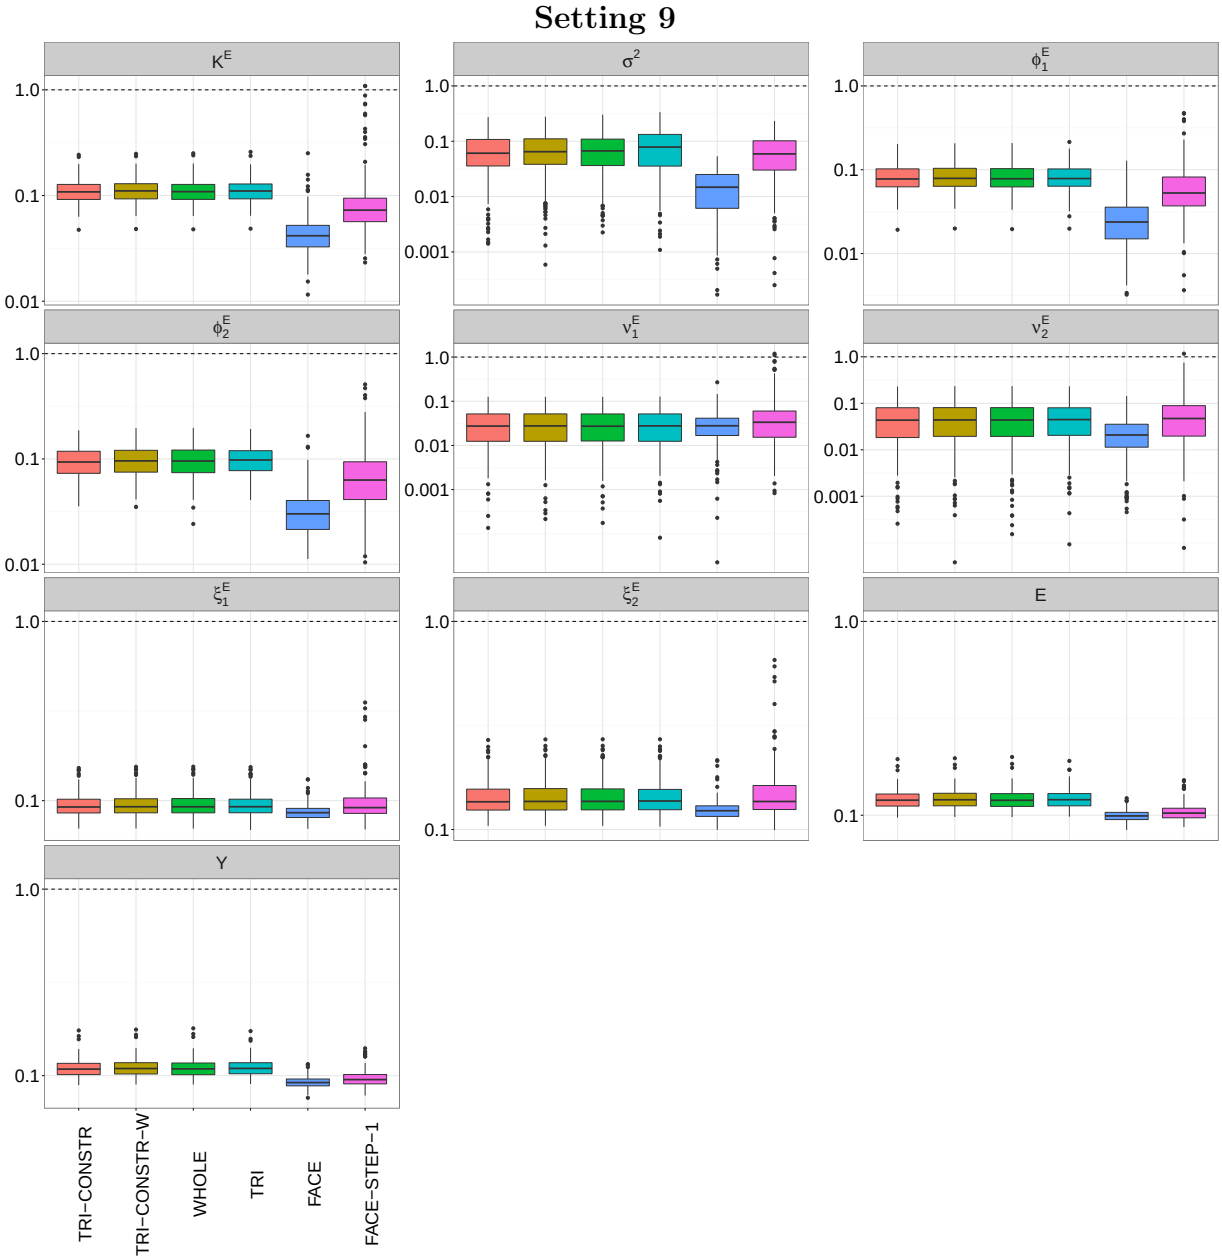

Figure 19: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 9 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

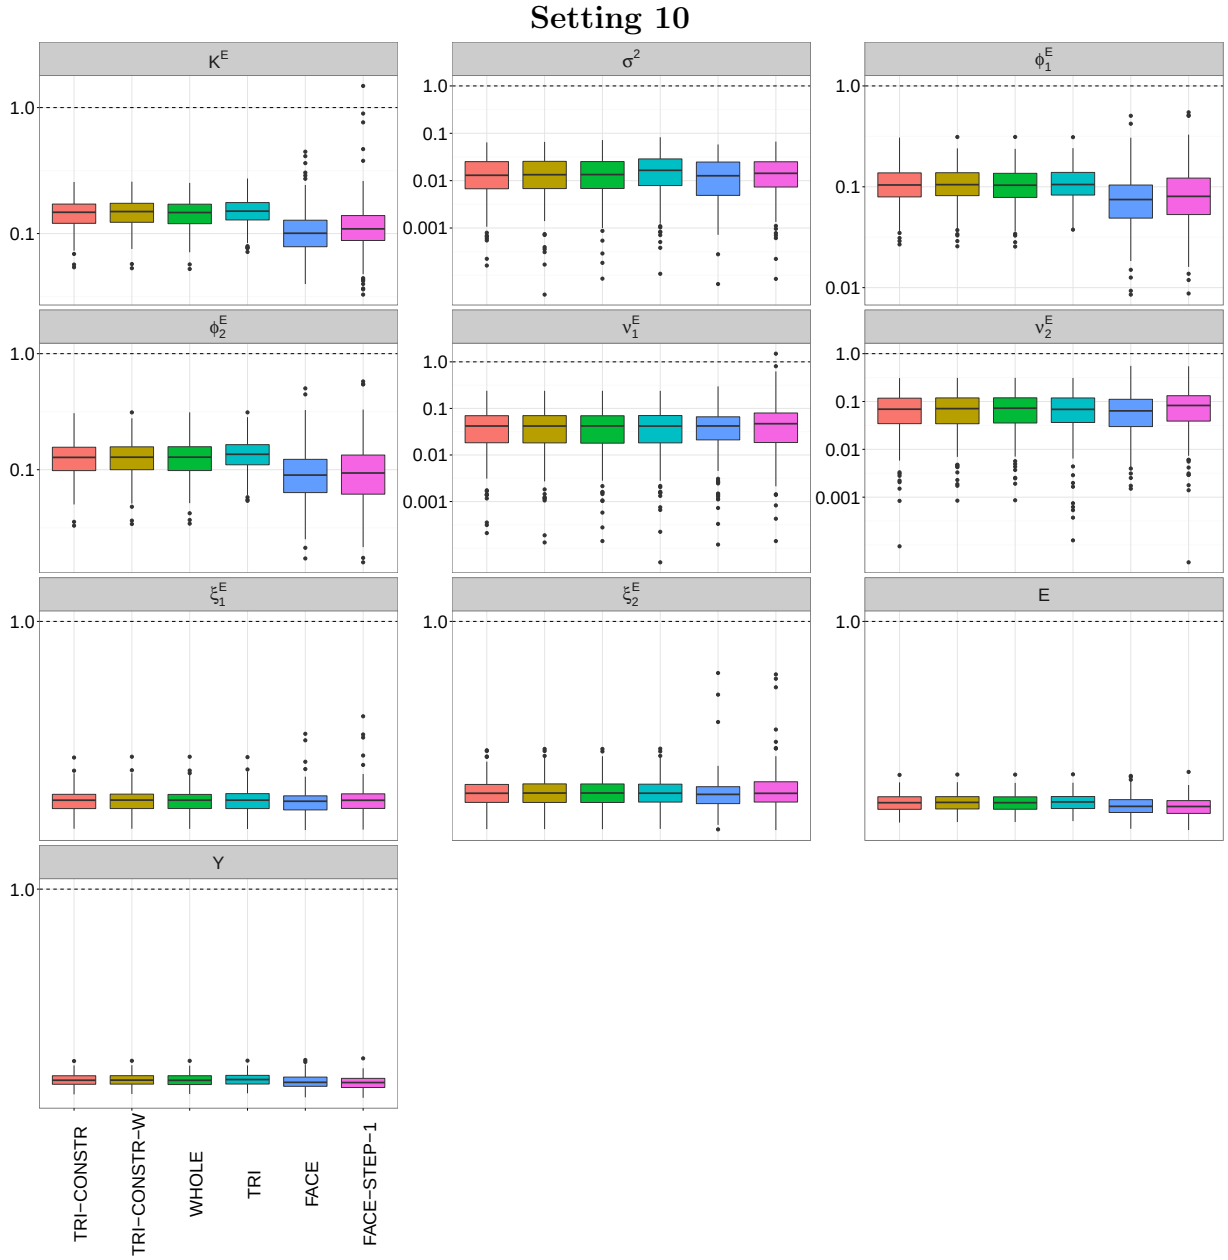

Figure 20: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 10 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

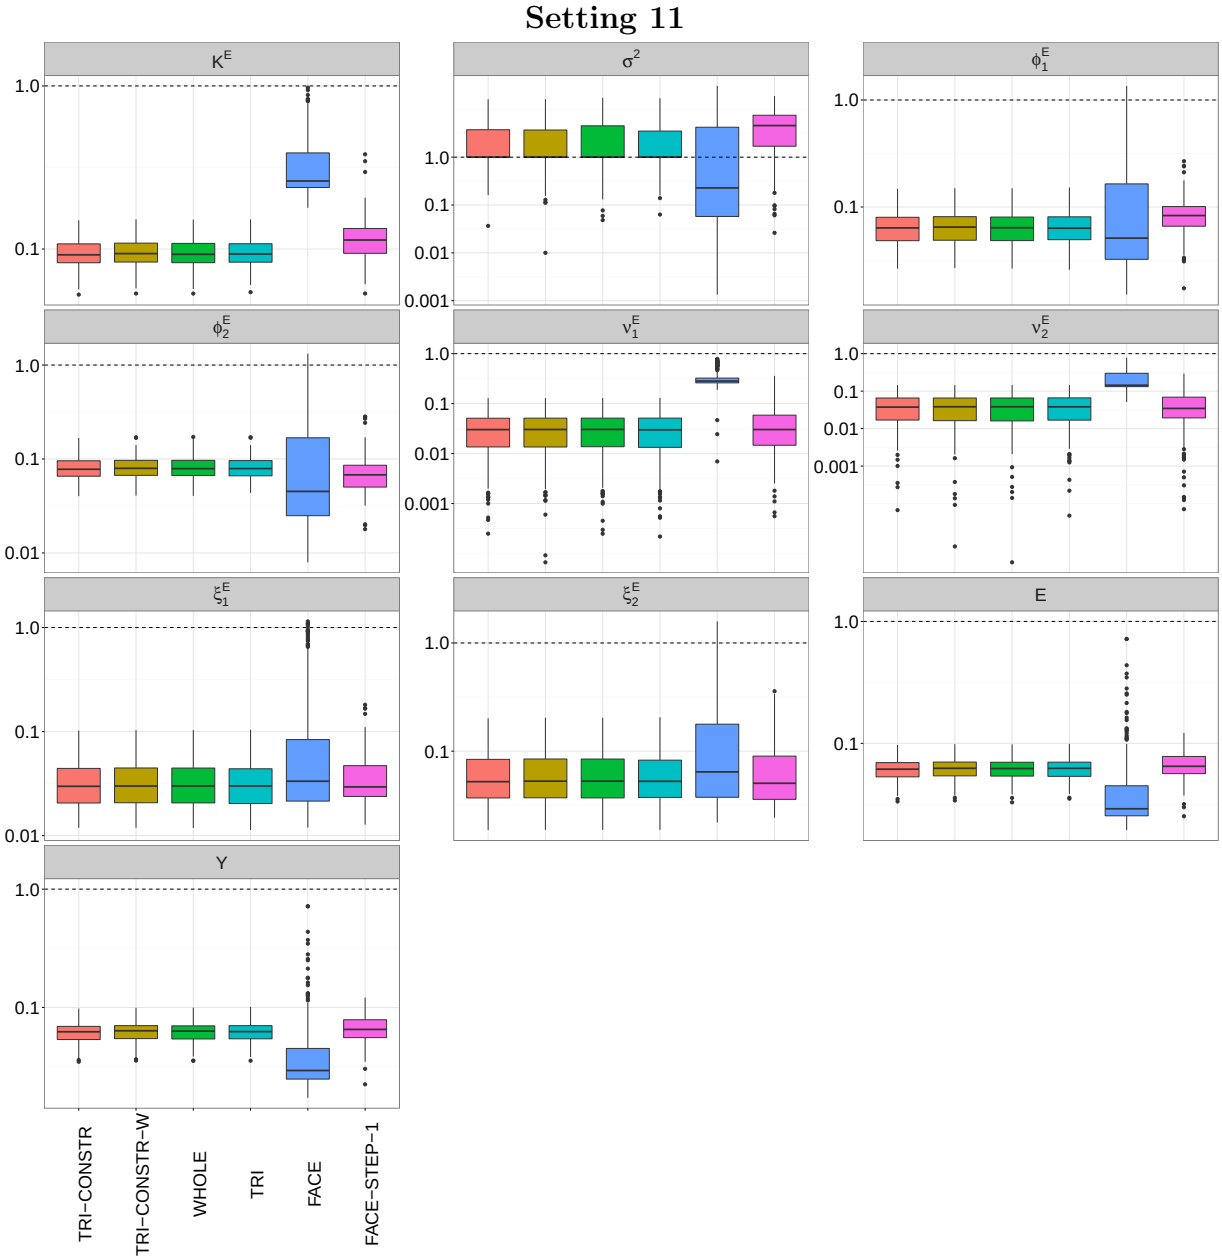

Figure 21: Boxplots of the rrMSEs (log10 scale at y-axis) for Setting 11 of the scenario with independent curves. Top row: rrMSEs for auto-covariance $K^E(t,t')$, error variance σ^2, and the first eigenfunction $\phi_1^E(t)$. Second row: rrMSEs for the second eigenfunction $\phi_2^E(t)$ and eigenvalues ν_1^E, ν_2^E. Third row: rrMSEs for the random basis weights ξ_1^E, ξ_2^E and process $E_i(t)$. Bottom row: rrMSEs for curves $Y_i(t)$.

D.4 Results for the scenario with crossed fRIs

In the following, we show the remaining results for the scenario with crossed fRIs. Figure 22, shows boxplots of the rrMSEs for the estimated eigenfunctions and eigenvalues, as well as for the random basis weights for the three random processes $B_i(t)$, $C_i(t)$, and $E_i(t)$. All boxplots are based on 200 simulation runs.

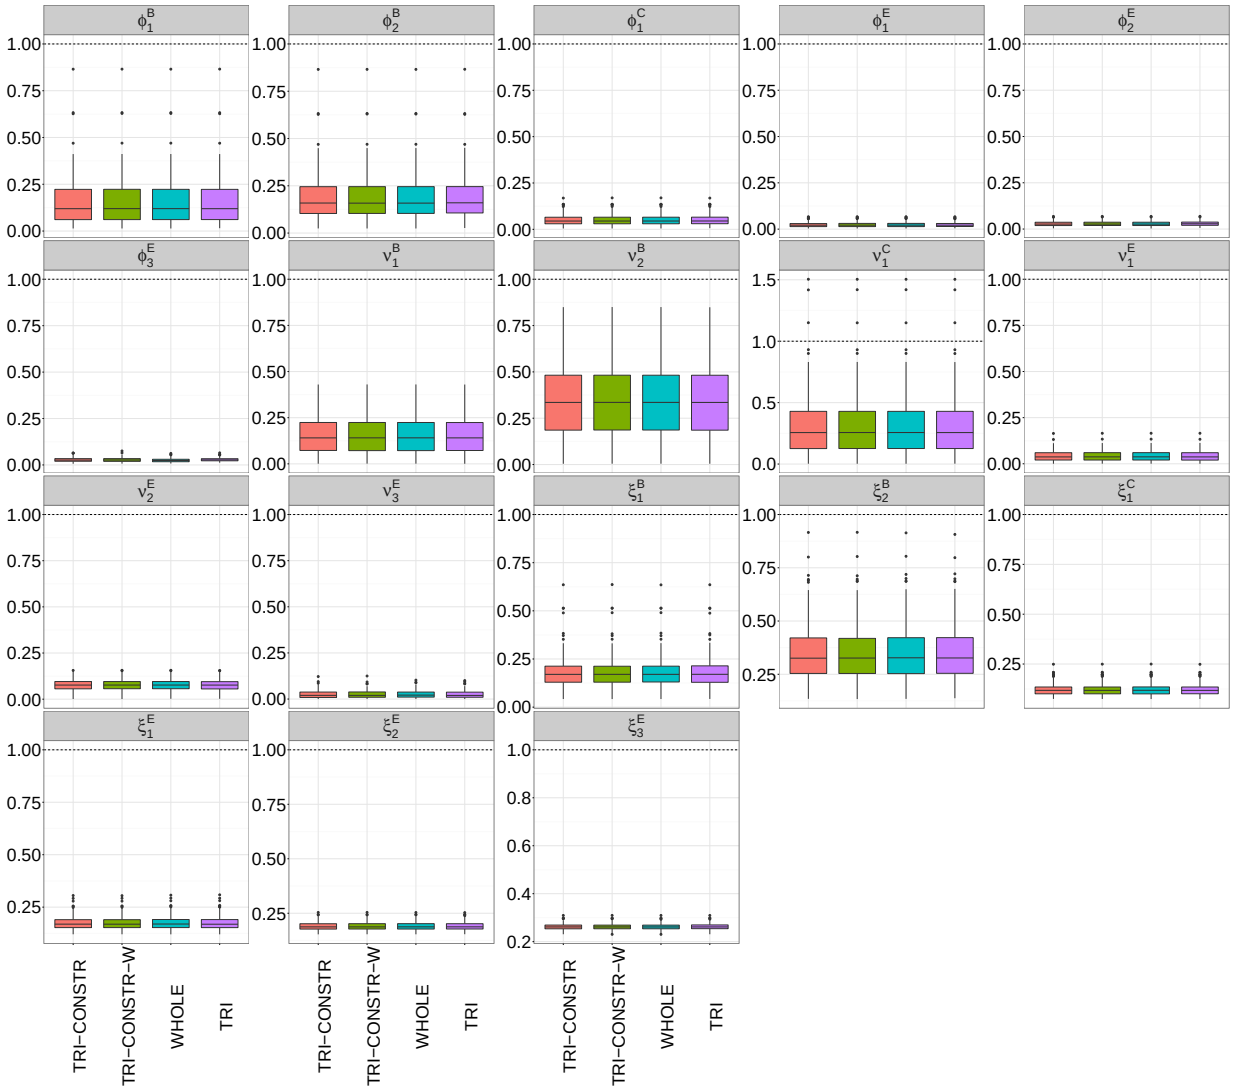

Figure 22: Boxplots of the rrMSEs for the crossed fRIs setting. Shown are boxplots for all remaining model components, which are not shown in Section 7.2: The rrMSEs for the eigenfunctions and eigenvalues of the three auto-covariances $K^B(t,t')$, $K^C(t,t')$, and $K^E(t,t')$, as well as the rrMSEs for the corresponding random basis weights.